DIGITAL POWER WEEKLY

May 4 – May 10, 2026

Editor: *Chris Ruppel, SVP - Power*

ISSUE NO. 1
KEEL INFRASTRUCTURE

$30B	300 MW	$9.8B	63GW
Anthropic Annualized Revenue Run-Rate (vs. $9B EOY 2025)	Anthropic Lease of SpaceX Colossus 1 (220K Nvidia GPUs)	Hut 8/Beacon Point 15-yr Lease (1 GW ERCOT, IG tenant)	AEP Incremental Contracted Large Load Expected By 2030

EXECUTIVE SUMMARY

The week's defining story was Nvidia's first extension of its $40B+ strategic investment campaign beyond the silicon and optical sectors into the land-and-power layer of the AI stack. Through a $2.1B equity option in IREN—paired with a $3.4B AI Cloud contract and a 5 GW DSX deployment commitment—Nvidia effectively signaled that Digital Power is now part of the strategic AI infrastructure stack.

Compute Scarcity = Power Scarcity. Anthropic leased 300 MW and 220,000+ Nvidia GPUs from SpaceX's Colossus 1 — capacity xAI had been running at ~11% utilization — rather than wait on conventional cloud, greenfield campus, or utility interconnection timelines (terms were not disclosed). While Anthropic's annualized revenue run-rate reached ~$30B (vs. ~$9B EOY 2025), it was facing documented usage limits, reliability complaints, and peak-hour frustration evidencing demand outrunning supply. The risk of losing market share was palpable.

Hyperscale Project Finance Expands. Hut 8/Beacon Point signed a 15-year, 352 MW IT lease (1 GW campus ceiling) at $9.8B base/up to $25.1B with extensions to an investment grade tenant in Nueces County, TX; shares had the largest single-day move since 2021. Meta's $13B El Paso financing mandate to Morgan Stanley/JPMorgan extends Louisians's Project Beignet ($30B, $27B debt) template into repeatable project-finance form at hyperscale-campus scale.

Utility Capex Grows/Regulatory Turmoil Escalates. IOUs plan ~$1.4T of capex over five years (+20% YoY; ~matches the entire 2015–2024 decade at $1.3T). PJM utility CEOs (AEP, FirstEnergy, Exelon) used the cycle to criticize current market structure; PJM CEO David Mills publicly called the design "not tenable." AEP raised its capital plan to $78B against 63 GW of expected incremental load by 2030 (~90% data-center-linked), floated a PJM/SPP exit, and signaled a West Virginia non-utility GenCo. MISO summer capacity cleared at $424.30/MW-day in North/Central. PA PUC adopted the model large-load tariff 5-0 on April 30.

Storage Switches to AI Infra. Fluence Q2 revenue missed consensus by ~24% ($465M vs. ~$615M), but backlog reached a record $5.6B with two hyperscaler MSAs disclosed and FY26 guidance ($3.2–3.6B) reaffirmed. Shares roughly doubled off pre-print levels on a 5-session run. BESS is being underwritten as interconnection and ramping infrastructure for AI campuses, not only as a renewable-attached asset.

BTC-to-AI Platforms Accelerate. Q1 prints documented the HPC pivot despite GAAP misses driven by infrastructure build-out and BTC mark-to-market: IREN $144.8M; WULF $34M total / $21M HPC (+117% sequential, 60 MW Core42 energized at Lake Mariner); HUT $71M (+226% YoY, $16.8B aggregate signed lease portfolio); CIFR $34.8M ($200M revolver closed, third hyperscale lease signed). The cohort is being repriced as contracted digital-infrastructure platforms.

★ MARKET SPOTLIGHT | FEATURED ANALYSIS

NVIDIA $40B+ EQUITY SPREE — IREN $2.1B OPTION/$3.4B AI CLOUD/5 GW DSX, CORNING $500M WARRANT PACKAGE/UP TO ~$3.2B EQUITY EXPOSURE

Sources: Bloomberg / Lynn Doan, Dina Bass (May 7, 2026); Bloomberg / Lynn Doan (May 6, 2026); CNBC / Jordan Novet (May 9, 2026); NVIDIA & IREN press releases / GlobeNewswire (May 7, 2026); IREN 8-K filing (May 7, 2026); Reuters (May 7, 2026)

Nvidia's $40B+ strategic-investment, warrant, prepayment, and commercial-capacity campaign now spans optics, neocloud, and powered AI infrastructure: Coherent and Lumentum at $2B each, Nebius at $2B, Corning up to $3.2B, and IREN up to $2.1B — before considering broader reported commitments such as Intel and OpenAI. The Intel position alone has compounded from a $5B bet to >$25B since announcement, which is why the playbook is now industrialized. Several of the public-company deals use warrant structures struck above spot, while other parts of the campaign involve direct equity, prepayments, or commercial capacity commitments. Corning: 15M warrants at $180 vs. $162.10 close. IREN: 30M at $70 vs. $56.85 close. Lumentum and Coherent in March were described publicly as $2B investments with purchase commitments/capacity access. Nvidia gets inexpensive optionality; the partner gets the announcement halo and a balance-sheet validation event.

- **Corning (May 6, $500M Nvidia-linked warrant package/up to ~$3.2B potential equity exposure + optical capacity buildout.).** Nvidia committed $500M in connection with Corning's optical manufacturing expansion and received a pre-funded warrant for up to 3M shares plus a five-year warrant to purchase up to 15M shares at $180 per share. If fully exercised, Nvidia's total potential equity exposure would reach approximately $3.2B. Corning shares jumped 12% in the regular session and as much as 18% in pre-market the next day; the stock has now compounded over 300% trailing twelve months. Three new optical-manufacturing complexes will be built in North Carolina and Texas, creating at least 3,000 jobs and increasing Corning's U.S. optical fiber production capacity by more than 50% (with optical connectivity capacity rising 10x). CEO Wendell Weeks separately disclosed to CNBC's Jim Cramer that Nvidia is funding factory construction outright in addition to the equity option — Nvidia has prepaid significant amounts for Corning's CapEx, an unusual inversion of the customer-vendor relationship. The strategic logic: copper interconnects inside Nvidia's Vera Rubin and Rubin Ultra rack-scale designs are being replaced by co-packaged optics, and Corning's glass-core packaging plus low-loss optical fiber are the binding upstream input. The Corning print stacks on top of the $4B Coherent + Lumentum equity options Nvidia struck in early March (laser/optical components), forming a coherent multi-billion vertical play across the photonics layer.
- **IREN (May 7, $2.1B equity option + $3.4B AI Cloud contract + 5 GW DSX deployment).** Per the joint press release and IREN's same-day 8-K, IREN issued Nvidia a five-year right to purchase up to 30M ordinary shares at a $70 exercise price (vs. $56.85 close on May 7, ~23% above spot) — equity participation up to $2.1B. The equity is paired with a separate $3.4B AI Cloud contract for Blackwell-generation GPU capacity and a strategic partnership to deploy up to 5 GW of NVIDIA DSX-aligned AI infrastructure across IREN's global pipeline. The flagship deployment path centers on IREN's 2 GW Sweetwater platform in West Texas, including Sweetwater 1 in Fisher County. IREN shares rallied on the Nvidia announcement and traded above the May 7 close the following day, before later volatility tied to financing needs for the AI buildout. IREN ended April with $2.6B of cash and previously signed a $9.7B AI compute contract with Microsoft (November 2025). Concurrent with the Nvidia announcement, IREN disclosed the acquisition of Ingenostrum SL, a Spanish data center developer, to support international expansion.

Stack mapping. Lay the six counterparties on the AI infrastructure stack and the picture emerges: **(i) optical/photonics layer** — Corning ($3.2B, fiber + glass-core packaging), Lumentum ($2B, lasers), Coherent ($2B, optical components); **(ii) networking silicon** — Marvell ($2B, custom ASIC + co-packaged optics, defensive against Broadcom); **(iii) neocloud / GPU-as-a-service** — Nebius ($2B, alongside prior commercial commitments to CoreWeave); **(iv) land + power + data center + GPU operations (vertically integrated)** — IREN ($2.1B equity option + $3.4B AI Cloud + 5 GW DSX commitment). IREN provided the upstream power and grid layer. Previously, Nvidia had been willing to backstop the silicon, the optics, and

the cloud reseller — the IREN deal moves Nvidia closer to the powered-land layer, though it remains an equity-option/commercial-capacity structure rather than a direct investment in generation or transmission. IREN closed at $56.85 on May 7; the $70 strike implies ~23% pre-commit upside.

Read-through. Nvidia's strategic-investment campaign now spans every layer of the AI factory: photonics (Corning, Lumentum, Coherent), networking silicon (Marvell), neocloud (Nebius), and — with IREN — digital power for the first time. The other five counterparties are vendors to the building; IREN is the building.

TOP STORIES

HUT 8/BEACON POINT — A $9.8B LEASE TURNS ERCOT POWER ACCESS INTO CONTRACTED AI REVENUE

Sources: Bloomberg / Dina Bass, "Hut 8 to Lease Texas AI Data Center for at Least $9.8 Billion" and "Hut 8 Jumps Most Since 2021 on Texas AI Data Center Lease" (May 6, 2026).

Hut 8 signed a 15-year lease to provide computing power to an investment-grade tenant at Beacon Point in Nueces County, Texas, valued at $9.8B base, up to $25.1B with three five-year extension options.

Phase 1 leases 352 MW of IT capacity at a 1 GW campus, with AEP Texas as utility and first energy delivery targeted Q1 2027. Hut 8's aggregate signed lease portfolio across Beacon Point and River Bend (LA) is $16.8B. Jacobs is EPC; Vertiv supplies cooling. Hut 8 closed with the largest single-day gain since 2021; Q1 revenue $71M (+226% YoY).

IREN/SWEETWATER/NOSTRUM — 2 GW IN TEXAS, 490 MW IN SPAIN, AND THE GLOBALIZATION OF THE MINER-TO-AI MODEL

Sources: Data Center Dynamics / Dan Swinhoe, "Iren energizes first phase of 2GW Sweetwater data center campus in Texas" (May 6, 2026) and "Iren acquires Spanish data center developer Nostrum" (May 8, 2026); Bloomberg / Lynn Doan and Dina Bass, Nvidia / IREN partnership (May 7, 2026).

Two execution updates accompany the Nvidia structure. IREN energized the high-voltage substation at Sweetwater 1 (1.4 GW designed campus, Fisher County, TX), with the initial 200 MW IT phase under construction. Sweetwater 1 sits on 1,700 acres; Sweetwater 2 (500 acres) targets 2028 energization. The combined two-site campus is 2 GW across ~2,200 acres. IREN reports >23,000 GPUs in fleet with ~11,000 under contract, Microsoft identified as a major customer.

Separately, IREN agreed to acquire Spanish developer Ingenostrum / Nostrum Group, adding ~490 MW of secured grid-connected power in Spain across Extremadura, Castilla-La Mancha, and Galicia plus a ~50-person development team. The acquired division had also claimed 5 GW of renewable capacity live and another 1 GW in development, adding a power-development channel to IREN's AI cloud expansion.

Read-through: The miner-to-AI model is globalizing. Sweetwater proves IREN can energize at GW campus scale in ERCOT; Nostrum gives it a European grid-connected pipeline and a local development organization. The acquisition also answers the obvious question after the Nvidia deal: how does IREN keep feeding a 5 GW DSX-aligned deployment path? It buys development capability and power positions, not just GPUs.

ANTHROPIC/SPACEX — 300 MW OF COLOSSUS CAPACITY

Sources: Wall Street Journal / Elias Schisgall (May 6, 2026); Bloomberg / Shirin Ghaffary and Carmen Arroyo (May 6, 2026); New York Times / Natallie Rocha (May 6, 2026); Bloomberg Opinion / Dave Lee (May 7, 2026); Bloomberg Businessweek / Max Chafkin, Musk v. OpenAI courtroom drama (May 7, 2026).

Anthropic leased Colossus 1 capacity from SpaceX, in scale and structure unusual for a frontier-lab tenant.

Anthropic agreed to use computing capacity at SpaceX's Colossus 1 data center near Memphis. The agreement gives Anthropic access to more than 300 MW of computing capacity and more than 220,000 Nvidia GPUs by the end of the month. Financial terms were not disclosed, but Bloomberg noted that recent data-center tenant economics have ranged from roughly $1.5M to $2.0M per MW per year, implying that a 300 MW commitment could represent hundreds of millions of dollars of annualized business for SpaceX / xAI.

Demand backdrop: At Anthropic's recent developer conference, Dario Amodei said the company had planned for ~10× growth this year and is tracking faster potentially reaching 80x. Anthropic said the SpaceX capacity will allow it to raise usage limits across Claude products, including Claude Code and the Claude API.

The capacity pinch is also becoming a competitive vulnerability. Bloomberg Opinion highlighted rising client frustration around Claude reliability, including Jefferies checks showing Claude Code uptime of 99.18% versus 99.98% for OpenAI's Codex, with the gap attributed to token constraints, compute limitations, and technical debt. Anthropic acknowledged that unprecedented demand had stretched its infrastructure, particularly at peak hours.

Supply side: xAI / SpaceX had been stockpiling Nvidia chips and building data centers ahead of model demand. An internal memo cited by Bloomberg put xAI chip utilization at ~11%. Leasing Colossus 1 to Anthropic converts excess to merchant compute revenue at scale.

Read-through. A frontier lab chose to lease 300 MW from a competitor rather than wait on utility queue, greenfield campus, or hyperscaler-region buildouts. Speed-to-compute is now coupled to speed-to-power.

AI labs are now paying strategic premiums for delivered capacity. Frontier labs appear willing to lease tenant capacity from competitors when megawatts are deliverable, which positions power-first operators as merchant suppliers into the frontier-lab demand stack.

CORE SCIENTIFIC/POLARIS — 440 MW OF OKLAHOMA INTERCONNECTION

Sources: Energy.Media, "Core Scientific Acquires Polaris Technologies & 440 MW Grid Connection" (May 7, 2026).

Core Scientific agreed to acquire Polaris DS LLC for approximately $421M in cash, with up to $40M of contingent consideration tied to incremental firm capacity. The transaction gives Core Scientific contracted access to 440 MW of gross power under an OG&E energy agreement, plus land, on-site substation infrastructure, and development rights at its Muskogee, Oklahoma campus. Core Scientific is already building a 70 MW facility for CoreWeave near the same campus and is targeting 1.5 GW of gross power through a combination of grid-connected and behind-the-meter solutions. The headline purchase price implies roughly ~$950–$1,000/kW of contracted gross power before allocating value to land, substation infrastructure, development progress, or contingent capacity.

PRIME/UGI — PENNSYLVANIA GAS-TO-POWER

Sources: Data Center Dynamics / Zachary Skidmore, "Prime acquires land parcel in Pennsylvania with natural gas supply from UGI" (May 8, 2026).

Prime's Pennsylvania transaction is undisclosed in financial terms but large in strategic signal. UGI Energy Services and Prime signed a northern Pennsylvania natural-gas infrastructure agreement under which UGI will sell Prime property for a proposed on-site gas-fired generation facility. UGI will retain approximately 15 Bcf of underground storage capacity and associated oil and gas rights tied to the property, while the site benefits from nearby Marcellus gas, interstate pipeline systems including Eastern Gas Transmission and Storage and Tennessee Gas Pipeline, and UGI Utilities distribution infrastructure. This is the Pennsylvania version of speed-to-power: not merely land, but land plus gas supply adjacency, pipeline optionality, utility coordination, and a power-generation pathway.

Read-through: Pennsylvania's competitive angle is no longer just PJM access. It is Marcellus gas, brownfield industrial land, storage, pipeline optionality, large-load tariff reform, and proximity to East Coast compute demand centers.

UTILITY ACTIVITY WRAP-UP — THE REGULATED UTILITY MODEL HITS THE DATA CENTER WALL

Sources: AEP Q1 2026 earnings; Bloomberg / John Worrall, May 8, 2026 (AEP PJM/SPP exit threat); PJM CEO David Mills letter and policy statement, May 2026; Virginia State Corporation Commission Dominion large-load interconnection hearing transcript (3 days, 1,126 pp); MISO 2026/2027 Planning Resource Auction results (April 2026); Keel Infrastructure strategic memorandum on PJM Powering Reliability Through Market Design (May 7, 2026); PA PUC Model Large Load Tariff final order summary (April 30, 2026); Data Center Dynamics / Zachary Skidmore, Prime / UGI Pennsylvania gas supply transaction (May 8, 2026).

Utility prints this week point to a shift: the large-load / data center problem has moved from a "load growth" story to a market-design, cost-allocation, and speed-to-power crisis. AEP, Dominion, PJM, and MISO are all now wrestling with the same core question: how do regulated utilities serve hyperscale load fast enough to win economic development, while avoiding stranded-cost risk, ratepayer backlash, and reliability deterioration?

UTILITIES ARE NO LONGER TREATING DATA CENTER DEMAND AS NORMAL LOAD GROWTH

AEP's Q1 update: AEP now expects 63 GW of incremental load by 2030, backed by signed large-load agreements, with roughly 90% tied to data centers, up from 56 GW the prior quarter. The load is concentrated across 41 GW in ERCOT, 16 GW in PJM, and 6 GW in SPP, with a much larger 190 GW active interconnection pipeline behind it. AEP's capital plan was raised from $72B to $78B, driven by new PJM/SPP transmission projects and gas-fired generation investment, while management also pointed to another $10B of potential projects not yet in the base plan.

AEP frames data centers as sources of fixed-cost contributions, citing up to $16B of "offsets" for existing customers if large loads are properly contracted. That is the structural argument utilities will use against the "data centers raise rates" political claim.

AEP'S PJM/SPP EXIT THREAT IS POSTURE — BUT THE POSTURE MATTERS

AEP's threat to evaluate leaving PJM and SPP should not be read literally. AEP is not likely to exit PJM in any practical near-term sense. The company is too embedded in PJM/SPP transmission, has billions of dollars of newly approved projects, and would face a legally and politically brutal separation process. Bloomberg notes that leaving PJM would be costly, difficult, and would require federal regulatory approval. But the threat is still meaningful. CEO Bill Fehrman's message was that utilities cannot wait a decade for RTO stakeholder processes to solve the generation-interconnection and speed-to-market problem. His line that, without action, *"we could still be having these same conversations in 10 years"* is the real takeaway.

The more important signal is not "AEP may leave PJM." AEP is reportedly considering participation in PJM's reliability-backstop construct through an unregulated generation company and a potential West Virginia non-utility GenCo model. More broadly, those comments point toward a hybrid market structure in which utilities remain inside RTOs while adjacent bilateral, GenCo, BYOC, or BTM structures develop to serve large loads faster than the standard queue process. The NIPSCO/Amazon CPCN-waiver filing in Indiana (Issue No. 6) is the same architecture from a different vantage point.

PJM IS NOW OPENLY ADMITTING THE CURRENT DESIGN IS NOT TENABLE

PJM CEO David Mills' May 6 letter and the accompanying PJM policy framing (white paper entitled *Powering Reliability Through Market Design*) are the institutional counterpart to AEP's complaint. Mills' statement that the *"current situation is not tenable"* matters because it acknowledges that PJM's problem is deeper than a temporary capacity-price spike or a narrow interconnection backlog. PJM is trying to balance three goals that are now in direct conflict: attracting new supply, protecting consumers from rate shock, and maintaining reliability during an AI-driven load surge.

"Credibility Trap." PJM's paper describes what is effectively a capacity-market credibility trap: scarcity prices are supposed to attract new supply, but if those prices trigger affordability interventions, investors discount the durability of future revenues and capital stays on the sidelines. This creates a "disequilibrium" where capital stays on the sidelines because the risk of policy reversal is perceived as higher than the potential return, leaving the grid with a growing generation deficit even as prices remain high. PJM's emerging answer appears to be some combination of differentiated reliability, backstop procurement, bilateral contracting, and large-load obligations — essentially, new rules that make hyperscale load bring money,

capacity, flexibility, or all three. The Reliability Backstop Procurement (RBP) auction PJM previewed on April 10 (Issue No. 6) is a concrete step to addressing the 'Credibility Trap.'

Mills framed the timeline as "years, not decades" for redesign; AEP's Fehrman said the speed-to-market problem "clearly has to get fixed." The two together document a structural capacity-market credibility gap and institutional escalation from AEP against an RTO process that is not moving at hyperscaler speed.

DOMINION SHOWS THE SAME FIGHT AT THE STATE-REGULATORY LEVEL

The recent Virginia proceeding (May 5) is the state-level version of the same conflict. The hearing on Dominion's large-load interconnection standards involved a who's-who of data center and infrastructure players — **Emerald AI, Google, Verrus, Tract, Amazon, and others** — and generated three days and 1,126 transcript pages. The dispute areas are exactly what you would expect in the new data center power regime: (1) whether interconnection standards can be separated from ratemaking; (2) whether large-load customers should make stronger financial commitments; (3) whether speculative load should be filtered out of forecasts; (4) whether transmission and substation costs should be directly charged to data centers; and (5) whether customers with executed Electric Service Agreements (ESAs) should receive queue priority.

This is important because Dominion is not arguing about whether the load is coming. The fight is over **who pays, who gets priority, and who bears stranded-asset risk**. This could become a template for large-load proceedings in other states in which queue access will increasingly be tied to financial deposits, contract maturity, ramp schedules, environmental compliance, and some showing that the load is real rather than speculative. The DTE Energy contract architecture for Oracle and Google in Michigan (Issue No. 6) is a similar playbook, executed.

MISO'S CAPACITY AUCTION SHOWS THE RELIABILITY-PRICE SIGNAL SPREADING BEYOND PJM

The MISO 2026/2027 Planning Resource Auction adds another regional datapoint. Fall, winter, and spring cleared at relatively low levels — **$33.92, $35.97, and $7.61/MW-day** — but summer prices cleared much higher, with **$424.30/MW-day in the North/Central zones**. MISO attributed the outcome to its Reliability Based Demand Curve, which is designed to price capacity based on reliability value above and below reserve targets. That matters because the capacity market is no longer just a PJM story. Across regions, the market is beginning to price the scarcity value of dispatchable, accredited capacity. Even where annual averages look manageable, **summer reliability is becoming the binding constraint**. For data centers, that means "energy price" is increasingly the wrong metric. The real cost stack is energy plus capacity plus transmission plus upgrade collateral plus curtailment / flexibility obligations.

POWERING RELIABILITY THROUGH MARKET DESIGN— PJM WHITE PAPER AND PA PUC NOW PUT THE "SUPPLY-BACKED LOAD" THESIS INTO FORMAL POLICY LANGUAGE

PJM is no longer merely saying capacity prices are high; it is saying high capacity prices alone do not create investment when investors lack revenue durability and regulatory stability over a financing horizon. That is precisely the gap that long-term bilateral offtake, backstop procurement, mandatory hedging, and supply-backed large-load structures are meant to close.

The recent white paper also provided key PJM data points that are now investment-grade talking points: the 2027/2028 BRA cleared at a 14.4% installed reserve margin versus a 20% target, implying a 6,517 MW UCAP deficit; empirical combined-cycle capital costs are routinely in the $2,000-$2,300/kW range for late-decade CODs; FID-to-COD timelines have doubled from roughly 24 months historically to 48+ months; and since 2020, roughly 24 GW of projects with executed interconnection agreements have terminated before COD, including 13.5 GW of natural gas projects.

PA PUC Large-Load Tariff Order (April 30) adds the state-level implementation layer. The Commission adopted the model tariff framework 5-0, with large-load thresholds of 50 MW individually or 100 MW aggregate at a single site, a but-for cost-causation standard, six-month EDC interconnection studies, public large-load queues, minimum five-year contracts, collateral requirements, early-termination cost recovery, universal-service contributions, and flexible / interruptible service options. The most important provision may be the self-construction pathway: large-load customers can build qualifying infrastructure upgrades, including broader network upgrades, subject to EDC safety, reliability, FERC, and NERC standards.

That combination matters for Pennsylvania. PJM is saying the market needs supply-backed, financeable megawatts; and the PUC is saying large loads must bear caused costs and may self-build infrastructure.

Read-through: The policy stack is moving toward exactly the architecture this report has been tracking. Large loads will be asked to bring money, collateral, flexibility, infrastructure, and in many cases supply. The winning development model is not "request XXX MW and wait." It is assemble site control, interconnection, gas, equipment, utility process, and tenant offtake into a financeable power-delivery product.

SUMMARY: UTILITY SECTOR'S EVOLVING POSTURE

- **Data center load is reshaping capital plans.** AEP's 63 GW contracted book and $78B capital plan document that utilities are underwriting growth around hyperscale demand, not studying it.
- **Standard RTO interconnection model is too slow for hyperscale timelines.** AEP's frustration with PJM/SPP and PJM's "not tenable" framing are expressions of the same institutional failure.
- **Large loads will increasingly have to prove seriousness with money.** Dominion's Virginia proceeding points toward deposits, ESAs, ramp discipline, direct assignment of upgrade costs, and transparency filings as the regulatory template.
- **Alternative structures are becoming mainstream.** The vocabulary now includes **GenCo structures, bilateral matching, reliability backstops, bring-your-own-capacity, flexible interconnection, and behind-the-meter or near-the-meter generation**..

Read-through. The question for the next 12 months is which utilities, states, and market structures can serve hyperscale load without breaking reliability or rate politics. Parallel structures (bilateral, GenCo, BTM, BYOC) are being built alongside the regulated tariff model rather than replacing it.

EARNINGS

AMD EARNINGS — INFERENCE DEMAND PULLS CPUs BACK INTO THE AI POWER STACK

Sources: Reuters / Zaheer Kachwala and Max A. Cherney (May 5, 2026); Bloomberg / Ian King (May 5, 2026); Wall Street Journal / Kelly Cloonan (May 5, 2026).

AMD's Q1 print extended AI infrastructure demand visibility beyond the GPU rack.

The company reported Q1 revenue of $10.25B, up 38% year-over-year, with adjusted EPS of $1.37. The data-center segment was the core driver: sales rose 57% to $5.8B, ahead of expectations, as EPYC server CPUs and Instinct GPU shipments scaled into hyperscaler AI demand. AMD guided Q2 revenue to $11.2B, plus or minus $300M, versus Street expectations around $10.5B, and the stock jumped roughly 12% in extended trading.

The more important number is not the quarter; it is the revised market framing. Lisa Su said AMD now expects the server CPU addressable market to grow at more than 35% annually and reach more than $120B by 2030 — roughly double the 18% annual growth rate the company had forecast in November. AMD also guided server CPU revenue to grow more than 70% year-over-year in Q2.

The shift from training to inference, agentic workflows, coding tools, and enterprise inference at scale moves more value toward CPUs (orchestration, memory movement, server density). AMD is positioning as the second-source accelerator against Nvidia and the share-taker against Intel in server CPUs.

There are constraints. Reuters and Bloomberg both highlighted the industrywide memory crunch: high-bandwidth memory demand for AI data centers is pulling memory manufacturers toward data-center products, pressuring PC component supply and pricing. AMD also remains dependent on TSMC manufacturing capacity, while Intel is attempting to reassert itself through in-house production.

Read-through. AI load growth is not GPU-only. Inference expands the compute surface to CPUs, memory, networking, and accelerators, raising rack densities and power envelopes regardless of chip mix.

AMD's data-center business is now the primary revenue and earnings driver. Guidance prices the CPU layer for AI. Inference demand pulls capacity toward distributed, latency-sensitive deployments closer to users, enterprises, and fiber routes, not only toward training clusters.

CUMMINS POWER — DATA CENTERS PULL BACKUP GENERATION INTO THE GROWTH STORY

Sources: Wall Street Journal / Bob Tita, "Cummins Power Business Takes Off, Thanks to China Data Centers" (May 5, 2026).

Cummins' earnings update extended the AI power cycle into generator-set demand and the industrial supply chain behind backup power.

Cummins raised its full-year revenue outlook to 8% to 11% growth, up from a prior 3% to 8% forecast, citing overseas data-center demand. Sales of power-generation equipment in China rose 84% in the first quarter from a year earlier, and the company now expects China revenue to increase 10% this year after previously forecasting a 1% decline. Cummins' power business revenue rose 19% in Q1, and the company noted it doubled production capacity last year for large engines used in generator sets.

Read-through. Generator availability joins turbine and transformer lead times as a supply constraint to track. Tier III / IV reliability standards keep diesel / gas gensets embedded regardless of hyperscaler procurement around fuel cells, batteries, and demand flexibility.

FINANCING

META EL PASO / SOPAIPILLA — HYPERSCALE CAMPUS FINANCE MOVES INTO PROJECT-DEBT SCALE

Sources: Bloomberg / Silas Brown, Aaron Weinman, Davide Scigliuzzo, Laura Benitez and Paula Seligson, "Meta Taps Morgan Stanley, JPMorgan for New Data Center Deal" (May 4, 2026).

Meta is reportedly working with Morgan Stanley and JPMorgan on a financing package for its El Paso data center that could total roughly $13B. The transaction is expected to be mostly debt with an equity component and follows the nearly $30B Beignet financing package for Meta's Louisiana campus, which included $27B of debt.

AI INFRASTRUCTURE CAPITAL MARKETS — REITS, GPU LOANS, HOT IPOS, AND $30B PLATFORM SALES

Sources: Bloomberg / Subrat Patnaik, Blackstone Digital Infrastructure Trust IPO (May 4, 2026); Bloomberg / Michelle Cheng and Jeannine Amodeo, CoreWeave GPU loan (May 5, 2026); Bloomberg / Anthony Hughes and Bailey Lipschultz, Cerebras IPO orders (May 5, 2026); Bloomberg / Elffie Chew, Pei Li and Manuel Baigorri, Stack Asia sale review (May 6, 2026); Bloomberg / Silla Brush and Layan Odeh, Larry Fink compute futures comments (May 5, 2026); Bloomberg / Silas Brown et al., Meta El Paso data center financing (May 4, 2026); Bloomberg / Echo Wong, Principal data center funds (May 7, 2026).

Four capital-markets prints extended the AI financing stack from private hyperscaler contracts into public REIT equity, syndicated GPU debt, chip IPOs, and platform M&A.

Blackstone Digital Infrastructure Trust is seeking to raise as much as $1.75B in a U.S. IPO, with shares expected to trade on the NYSE under BXDC. The vehicle targets newly built, leased data centers valued at $250M to $1.5B and leased to investment-grade hyperscalers, with Blackstone disclosing an estimated $1T stabilized data-center market opportunity over the next five years.

CoreWeave secured lower borrowing costs on a $3.1B GPU-backed loan after attracting more than $15B of investor orders. The margin tightened by 50 bps to 4.5 percentage points over the benchmark, with proceeds supporting GPU purchases and installations for OpenAI and Cohere contracts. The important structural point: GPU financing is moving from bespoke private credit into the syndicated leveraged-loan market.

Cerebras is seeing enough IPO demand that Morgan Stanley is requiring institutional buyers to submit limit orders specifying both share count and maximum price. Cerebras is seeking to raise as much as $4.8B (up from $3.5B a few weeks ago), The company is offering 30 million shares for $150 to $160 each, according to a filing with the US Securities and Exchange Commission on Monday, confirming an earlier Bloomberg News report. Cerebras had previously marketed 28 million shares for $115 to $125 each.

At the top of the price range, Sunnyvale, California-based Cerebras would have a market value of $34.4 billion, based on the shares outstanding in its filing. The IPO has drawn orders for more than 20 times the number of shares available, people familiar with the matter have said. The IPO is expected to trade on Nasdaq under CBRS.

Stack Infrastructure (Blue Owl-owned) is considering a sale of its Asia operations in Australia, Japan, and Malaysia, with a potential valuation above $30B. The review lands in a region already busy with data-center deal activity, including Digital Edge, Bridge Data Centres, Princeton Digital Group, and DayOne. Moody's expects at least $3T to flow into data-sector investment over the next five years, much of it debt-financed.

Read-through. The financing stack is catching up to the load curve: stabilized hyperscaler leases moving to REIT format, GPU contracts becoming loan collateral, chip supply bid like scarce infrastructure, regional platforms trading like strategic utilities. Capital is abundant for contracted digital infrastructure; the pricing premium will attach to assets that can document megawatts, interconnection status, and delivery timelines.

FERVO IPO

Sources: Bloomberg / Monique Mulima and Michelle Ma, "Geothermal Developer Fervo Seeks $1.33 Billion in US IPO" (May 4, 2026); Bloomberg Green / Michelle Ma and Naureen S. Malik, "Data Center Startup Fermi Ousts CEO After Stock Slides" (May 4, 2026).

Fervo is targeting up to $1.82B in a U.S. IPO at ~$7.37B implied market value (top of range). Fervo initially filed to raise up to approximately $1.33B during the May 4–10 issue window; subsequent May 11 reporting showed the IPO upsized to 70M shares at $25–$26, implying up to ~$1.82B raised and a valuation up to ~$7.37B.Cape Station in Utah is permitted as a 500 MW geothermal project, with permits for another 1.5 GW, ~2.6 GW in advanced development, and >38 GW in early-stage development. Fervo disclosed ~$7.2B of potential contracted PPA revenue and ~75% drilling-time reductions 2022–2025 (~70% drilling-cost reduction). The model adapts oil-and-gas drilling technique to 24/7 firm power.

SUSTAINABILITY

MICROSOFT 100/100/0 — HOURLY CLEAN POWER MEETS THE AI LOAD CURVE

Sources: Bloomberg */ Alastair Marsh and Brody Ford, "Microsoft in Talks to Ax Energy Pledge Amid Data Center Boom" (May 6, 2026).*

Microsoft is reviewing its 100/100/0 hourly clean-power target, set in 2021 to match 100% of electricity use, 100% of the time, with zero-carbon procurement from the same grids it draws from by 2030. Microsoft maintains its environmental commitment publicly, but the target is now under review. Annual renewable procurement is not the constraint. The constraint is hour-by-hour deliverability across constrained grids as AI campuses are added at gigawatt scale. Bloomberg reports Microsoft is adding ~1 GW of data-center capacity every three months.

Read-through: Much of the clean-power conversation has shifted from annual accounting to hourly operations. Microsoft's hourly clean-power framework is under pressure; the market should watch whether the company delays, narrows, or reinterprets its 100/100/0 target.

MISCELLANEOUS

GULF DATA TAKES THE PIPELINE ROUTE — FIBER RESILIENCE MOVES INTO THE ENERGY CORRIDOR

Sources: Rest of World / Indranil Ghosh, "Big Tech is moving data out of the Gulf through Iraqi oil pipelines" (May 1, 2026).

Hyperscalers are routing Gulf-region cloud traffic along Iraqi oil-pipeline corridors to reduce latency and create a land-based backup to vulnerable subsea cables.

Major U.S. hyperscalers operating in the Gulf-region data centers have reportedly bought capacity on an Iraqi route built by IQ Networks. The route runs from southern Iraq to the Turkish border and is now being extended through gas-pipeline corridors across Turkey to the European border, creating an eventual overland connection into Frankfurt, Amsterdam, London, and Marseille.

The demand for alternate routes intensified after Iranian drones reportedly struck Amazon facilities in the UAE and Bahrain, causing banking, payment, delivery, and enterprise-software disruptions across the region. Amazon told customers to migrate workloads out of the Middle East, while hyperscalers moved to diversify routes away from Red Sea and Strait of Hormuz subsea exposure.

Read-through. Energy corridors carry rights-of-way, security perimeters, access roads, and maintenance regimes that double as fiber-route infrastructure. Oil and gas pipeline corridors become latent platforms for latency arbitrage and geopolitical route redundancy alongside power and molecules.

G42 / CORE42 MINNEAPOLIS — SOVEREIGN AI MOVES INTO URBAN CONVERSION CAPACITY

Sources: Bloomberg / Natalie Wong, "UAE AI Firm Anchors New Downtown Minneapolis Data Center" (May 6, 2026).

Core42, a digital infrastructure unit under Abu Dhabi-backed G42, took **20 MW** at **1001 Third Avenue S.** in downtown Minneapolis, anchoring an office-to-data-center conversion. The site is part of Core42's global AI infrastructure buildout and adds another U.S. node to G42's footprint, which Bloomberg says already includes California, Texas and New York.

Read-through: This is the other side of the gigawatt-campus boom. Not every AI deployment will be a 1 GW exurban project. Sovereign-backed AI platforms, coding workloads, enterprise inference and latency-sensitive services will also seek smaller blocks of powered urban or near-urban capacity.

QII — CHAD WILLIAMS RETURNS WITH A GRID-ONLY BET ON NONTRADITIONAL MARKETS

Sources: QII market-entry analysis memorandum (May 2026).

QII combines Chad Williams' QTS track record, family-office capital, and senior energy-infrastructure talent in a platform built around powered shells, utility relationships, infratech supply-chain management, and non-traditional market selection. The strategy is grid-only: Williams is betting unutilized grid capacity exists in markets reachable via utility relationships, against the supply-backed BTM thesis that large-load grid service is structurally too slow in primary markets. The hire of former Siemens Energy NA president Rich Voorberg as Chief Infratech Officer signals that even the grid-only thesis is preparing for an eventual move toward owned or contracted generation.

GREAT AMERICAN SATELLITE AGE — SPACE CONNECTIVITY BECOMES PART OF THE DIGITAL INFRASTRUCTURE STACK

Sources: WIRED / Paresh Dave, "Welcome to the Great American Satellite Age" (May 4, 2026).

WIRED's recent satellite article points to the next communications-resilience layer around AI infrastructure. A new generation of satellite startups — Basalt Space, Xona, Earthtraq, Muon Space, Astranis, and others — is trying to make satellite imaging, positioning, navigation, timing, and communications more private, taskable, and resilient. The article frames this as a shift from government- and prime-contractor-dominated space infrastructure toward smaller, cheaper, software- and AI-enabled orbital networks.

Read-through is that "connectivity" is widening from fiber diversity into orbital redundancy. The same customers worried about undersea cables, Gulf routing, Starlink dependency, GPS spoofing, and sovereign data control are natural buyers of alternative satellite capacity. Astranis' reported $1B+ of customer commitments for country-specific satellite internet networks and Muon's wildfire-monitoring / remote-sensing use cases show that the digital infrastructure stack is no longer just data centers plus terrestrial fiber. For AI campuses, especially in remote power-rich locations, the long-term architecture likely includes redundant terrestrial fiber and satellite backup.

AWS US-EAST-1 OUTAGE — THERMAL EVENT FINANCIAL INFRASTRUCTURE RISK

Sources: Quartz / Cris Tolomia, "An AWS data center overheated, knocking out Coinbase for hours" (May 8, 2026).

The outage story this week was not a grid emergency, a cyberattack, or a software-only failure. It was a thermal event at an AWS data center in Northern Virginia that cascaded into a seven-hour Coinbase disruption.

Coinbase experienced a multi-hour disruption, reported by some outlets as roughly five to seven hours. The lesson is architectural. Even in a cloud region designed around multiple availability zones, a single data hall/power/cooling failure can still surface as a customer-facing outage when application architecture, data persistence, failover behavior, or dependency mapping is not sufficiently resilient.

Read-through. Thermal reliability is now a financial-infrastructure reliability input. Data-center underwriting will increasingly diligence cooling redundancy, UPS ride-through, thermal containment, and AZ-level dependency design as bankability inputs alongside power and tenant credit.

POWER, GRID & REGULATORY REPORTING CALENDAR — KEY DATES & UPCOMING CATALYSTS

Date	Regulatory / Reporting Catalyst	Why it matters for Digital Power
May 21, 2026	**FERC May Open Meeting / monthly order watch**	Monitor for large-load, PJM, co-location, transmission-planning, and eTariff orders that could shape data-center interconnection structures.
May 28, 2026	**FERC EIT decision requested watch (ER26-1563)**	Near-term ruling requested for the PJM Expedited Interconnection Track; defines whether qualifying generation / storage projects can jump the ordinary queue.
May/June 2026	**PJM 2028/2029 Base Residual Auction results window**	Watch auction execution and results timing as PJM returns toward its forward capacity schedule.
June 5, 2026	**PA DEP EIT RFI responses due**	Written responses due for Pennsylvania generation/storage projects seeking Commonwealth sponsorship into PJM EIT; responses reviewed on rolling basis.
June 2026	**FERC action expected in RM26-4-000 large-load ANOPR**	Highest-priority federal proceeding for significant-load interconnection; could set national guardrails for co-located load, cost allocation, transparency, and transmission-interface rules.
June 2026	**PJM Reliability Backstop Procurement (RBP) FERC filing / final design**	Watch for filing of the one-time backstop procurement framework: target MW, eligible resources, pay-as-bid/CfD structure, credit, collateral, term length, and cost allocation.
June 2026	**PA PUC model large-load tariff implementation watch**	Track EDC tariff/compliance activity following the April 30 model framework: thresholds, collateral, minimum term, termination liability, cost-causation, and self-build pathways.
June/July 2026	**Virginia SCC Dominion large-load queue case watch (PUR-2026-00011)**	State-level template for queue discipline: ESA maturity, projected-service dates, speculative-load filtering, collateral, minimum charges, and data-center priority rules.
Q2 2026	**PPL rate case/large-load tariff-class ruling expected**	Determines tariff treatment for large loads in PPL territory
Q3 2026	**WV PSC ruling on Mon Power 1.2 GW CCGT + 70 MW solar**	Regulated-utility generation precedent tied to PJM supply shortage; relevant to FE/WV structures and rate-base treatment of data-center-adjacent generation.
August 3, 2026	**NERC Level 3 Computational Load Alert reporting due**	Registered entities must report status of activities in response to NERC Essential Actions on modeling, studies, instrumentation, commissioning, operations, protection, and control for computational loads.
September 2026	**PJM RBP facilitated bilateral-contracting phase begins (proposal)**	PJM/CRA matchmaking phase for load, LSEs, large-load customers, generators, DR, and DER; key window for supply-backed load and bilateral capacity contracts.
September 2026	**EIA data-center energy-use pilot completion target**	Pilot survey output expected to inform a national data-center energy-use survey; metering, BTM generation, cooling, and server-energy reporting become bankability / compliance inputs.
Fall 2026	**Shippingport Lateral (FERC CP25-541) in-service target**	Gas-pipeline precedent for BTM/near-meter repowering; useful comp for Scrubgrass fuel-delivery and pipeline-permitting strategy.
October 2026	**Microsoft 100/100/0 clean-power target review**	Sustainability reporting catalyst: hourly matching, deliverability, firm power, storage, and grid congestion will matter more than annual REC accounting.
December 2026	**PJM 2029/2030 Base Residual Auction**	Second 2026 BRA; read-through for whether RBP, EIT, and large-load reforms are changing capacity-market confidence or merely repricing scarcity.
January 1, 2027	**Virginia GS-5 large-load safeguards begin for new contracts**	Dominion large-load customers contracting on/after this date face the new 14-year service obligation; minimum charges and collateral framework become active state precedent.
Q1 2027	**Hut 8 Beacon Point first energy delivery target (AEP Texas)**	Commercial validation point for contracted AI load served through utility-backed ERCOT delivery; compare against regulatory and interconnection milestones.
March 2027	**PJM RBP central procurement phase begins (proposal)**	If bilateral matching is insufficient, PJM central procurement opens; suppliers face feasibility, site control, equipment, financing, fuel, network-upgrade, and collateral gating.
By 2027 target	**First Air Force ANPI microreactors operational target**	NRC/federal deployment benchmark for advanced nuclear licensing and behind-the-fence power at mission-critical sites.