KEEL INFRASTRUCTURE

DIGITAL POWER WEEKLY

ISSUE NO. 7
June 15-21, 2026

EDITOR
Chris Ruppel
SVP - Power

$300B+
AI-LINKED DEBT ISSUANCE SINCE NOV.

$665B
OPENAI COMPUTE COMMITMENTS THROUGH 2030

$25B
NVIDIA BOND SALE; ~$85B ORDERS

6 RTOs
FERC SECTION 206 LARGE-LOAD ORDERS

In this Issue

WEEK IN REVIEW

AI Infrastructure Moves From Growth Theme to Industrial Stress Test

The key takeaway from this week was that the AI buildout is being repriced around execution: power access, financing capacity, regulatory tolerance, supply-chain credibility, cooling architecture and national-security treatment. The market is moving from "how many gigawatts were announced?" to "which projects can actually be energized, financed and defended?"

Capital remains available, but the numbers now look like infrastructure, not software. Nvidia sold $25B of high-grade bonds into roughly $85B of demand; SpaceX is preparing a $20B+ bond refinancing after its IPO; and OpenAI disclosed $3.7B of Q1 cash burn, $5.7B of revenue and $665B of compute commitments through 2030. The read-through is constructive but not complacent: the funding window is open, while the burden of proof on utilization, margins and power-backed execution is rising.

Regulators moved in the same direction. FERC's six-RTO large-load Orders to Show Cause, CAISO's computational-load ride-through proposal, Minnesota Power's Google demand-response structure, Virginia's data-center electricity tax and the DOJ/xAI Clean Air Act filing all point to the same conclusion: compute load is now grid policy, environmental policy, fiscal policy and national-security policy.

1. CAPITAL IS OPEN, BUT DISCIPLINE IS RISING

The market will increasingly separate durable, contracted, power-backed platforms from long-duration capex stories with unclear returns. Tenant credit, realistic schedules, upgrade responsibility and cash-conversion visibility matter more than announced capacity.

2. THE CANCELLATION NARRATIVE IS TOO BLUNT

SemiAnalysis' rebuttal to the "half of 2026 capacity is canceled" narrative is a useful correction. Delays are real, but many alarming estimates confuse speculative announcements with executable capacity. *SemiAnalysis* says its YE2026 North American hyperscaler self-build forecast moved only about 1% over six months, with North American colocation below 5%. The better framework is pipeline triage: real projects are moving—press-release GWs are discounted.

3. POWER IS NOW THE UNDERWRITING SCREEN

Development stories from Salem, Wyoming, Maryland, Meta/Crusoe, Google cooling retrofits and utility demand-response structures all point to the same rule: land is necessary but insufficient. The investable asset is controlled, operable power with a politically credible path to load.

4. OEMs AND BATTERIES ARE PIVOTING TO THE LOAD CENTER

Peak/GM, Schneider/Foxconn and Rolls-Royce SMR selection all reinforce the same supply-chain shift: energy equipment is being redesigned around the data-center customer, with storage, modular power, cooling and firm generation becoming part of the AI stack.

5. CHINA FRAMES THE COMPETITION

RAND's recent report *China's Techno-Industrial Strategy in the Xi Era: Producing Under Pressure* and *Bloomberg*'s articles point to the broader lesson: AI alone is not enough. The strategic race is manufacturing, power systems, supply chains, finance, industrial land, standards and deployment capacity. The U.S.'s advantage is capital formation and private innovation; the risk is underbuilding and/or delaying the physical systems needed to make the innovation matter.

BOTTOM LINE

This week AI infrastructure thesis became more investable by making the bottlenecks visible. Demand remains strong, but the market is learning to pay for execution rather than announcements. That favors platforms with real power, interconnection strategy, flexible-load architecture, credible financing and local political legitimacy. Leadership without energy execution and manufacturing scale is incomplete. The week's seemingly separate stories connect: FERC intervention, AI bond issuance, battery-OEM pivots, SMR orders, CAISO grid-code rules, Virginia taxation and Anthropic export controls are all symptoms of the same paradigm shift. AI is becoming primarily an infrastructure race, and infrastructure races are only won by countries—not companies.

Major links: Nvidia bonds | SpaceX AI fundraising | OpenAI cash burn | SemiAnalysis | Bloomberg FERC | Anthropic order | RAND China study | AI not enough

★ MARKET SPOTLIGHT | FEATURED ANALYSES

FERC'S LARGE-LOAD INTERVENTION: THE QUEUE FIGHT JUST BECAME FEDERAL POLICY

FERC's June 18 large-load package should be read as an intervention into the data-center bottleneck, not a narrow tariff cleanup. The Commission issued tailored Section 206 show-cause orders to PJM, MISO, SPP, CAISO, ISO-NE and NYISO, requiring each market and its transmission owners to defend current tariff treatment or propose revisions. The operating thesis is simple: faster load interconnection is now being paired with a tougher framework for cost causation, operational flexibility and resource adequacy.

WHAT FERC ACTUALLY DID

The legal posture matters. FERC has not imposed one national large-load tariff. It has forced a defend-or-revise process: each RTO/ISO must explain why existing tariffs are just and reasonable despite the absence of clear large-load provisions, or file tariff changes. A parallel 30-day report asks how each market intends to ensure adequate generation for existing and new large loads. The five required subject areas are transmission-service studies, cost-shift prevention, co-location and behind-the-meter generation, flexible large-load service, and study processes for proximate generation serving large or co-located load.

The public headline is speed - moving from multi-year hook-up uncertainty toward a 90-day paradigm. The more important commercial point is the trade: speed for responsibility. Large users may need to bring or contract for power, accept curtailment during system stress, and pay directly for upgrades. That turns the large-load debate from a simple interconnection queue problem into a full political-economy problem: who pays, who curtails, who gets priority and how much of the infrastructure burden is visible to ratepayers.

SPP IS THE REFERENCE ARCHITECTURE; PJM IS THE TEST CASE

SPP is the benchmark because its High Impact Large Load framework, paired-generation assessment and Load Limited Resource Interconnection Service concept show how new demand can be studied alongside proximate new supply. The strategic architecture is not just 'connect a data center faster.' It is 'connect a data center faster if the load can be paired with new power, limited service rights and operating conditions that protect the broader system.'
PJM is different because it is not starting from a blank page. Its co-location and behind-the-meter generation proceeding already includes interim network service, firm contract-demand service and non-firm contract-demand service concepts. FERC's next order cycle will determine whether PJM becomes a more financeable market for sophisticated, power-integrated data-center buyers - or simply a more transparent market with higher explicit costs and clearer curtailment exposure.

COMMERCIAL READ-THROUGH

The policy direction is constructive for serious projects and punitive for speculative ones. Land-only development becomes less valuable if it cannot answer the power question. Projects that can combine load, dedicated or proximate generation, flexible operations, fuel strategy, upgrade funding and tenant credit move up the stack. Pure load-only buyers that want immediate full-firm service with no upgrade exposure, no curtailment and no public cost transparency become harder to finance and harder to defend politically.

Read-Through: The market is moving from 'where can we plug in?' to 'who can show a controllable, financeable and politically defensible energization plan?' FERC is not eliminating regional tariff risk, but it is validating the power-integrated development model: bring power, pay for what you use, operate flexibly, and make the cost allocation visible.

Sources: FERC June 18 large-load integration release; FERC fact sheet; FERC June 2026 Commission meeting summaries; SPP HILL public materials; Bloomberg large-load coverage.

Links: FERC Large Load Integration | FERC Fact Sheet | FERC June 2026 Meeting Summaries | SPP HILL Integration | Bloomberg large-load coverage

6 RTOs/ISOs PJM, MISO, SPP, CAISO, ISO-NE, NYISO	**30 Days** Generation-adequacy reports	**60 Days** Defend-or-revise tariff response	**5 Buckets** Studies, costs, BTMG, flex load, proximate generation

Large-load intervention scorecard - the near-term regulatory calendar and commercial pressure points.

WHAT TO WATCH IN THE COMPLIANCE FILINGS

The market should not overread the 90-day framing until the actual tariff sheets land. The investable details will be mundane but decisive: whether "large load" is defined by MW threshold or system impact; how firm and non-firm contract demand are priced; whether behind-the-meter generation can net against load; how curtailment priority is assigned; whether upgrade costs are direct-assigned or socialized; and how existing or near-final service agreements are protected from rule changes.

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

WHY THIS MATTERS FOR DATA-CENTER FINANCING

The practical result is a cleaner underwriting screen. A project that can contract for capacity, document curtailment capability, stage load energization and accept direct upgrade exposure is now closer to the kind of customer RTOs can process. A project that only has land and a desired service date is increasingly the weak hand. The tariff answer may differ by market, but the commercial answer is converging: the more a customer can reduce system uncertainty, the more likely it is to receive a financeable path to service.

That has real capital-markets consequences. Debt and equity providers will start asking for evidence that a data-center developer/tenant is not just creditworthy, but grid-worthy: flexible enough to curtail, funded enough to pay upgrades, supplied enough to reduce resource-adequacy risk, and structured enough to survive tariff changes. FERC's intervention does not remove regulatory risk; it makes the risk legible.

THE 30/60-DAY CLOCK IS NOW THE MAIN WATCH ITEM

The next catalysts are procedural but important. The 30-day reports will show whether each RTO/ISO treats large load as a resource-adequacy problem, a transmission-service problem or both. The 60-day filings will show whether the markets are prepared to create new service products, new study queues, new cost-allocation rules and new operating obligations. The winners will not be the fastest press releases; they will be the first projects that can convert those tariff mechanics into bankable tenant commitments.

Dimension	SPP reference model	PJM test case
Vehicle	HILL load process, HILLGA paired-generation assessment and load-limited interconnection concepts.	E-2 co-location / BTMG track plus the new six-RTO Section 206 show-cause clock.
Timing signal	Publicly framed as faster, with paired-generation study architecture outside the legacy queue.	Still unsettled: further compliance response due within 60 days; final tariff treatment remains the gating item.
Economic bargain	New load becomes more acceptable when paired with new proximate supply and operating constraints.	Potentially more financeable if rules clarify minimum charges, netting, upgrade costs and curtailment priority.
Buyer implication	Best for sophisticated customers willing to bring power, fund upgrades and accept flexibility.	Potentially worse for load-only buyers seeking immediate full-firm service without system-cost exposure.

Bottom-line framework: *SPP is the public template; PJM is where the financing, pricing and curtailment details will determine whether the tariff becomes bankable.*

BANKABILITY CHECKLIST FOR LARGE-LOAD SPONSORS

The compliance filings should be read less as legal paperwork and more as the next underwriting package for AI load. The bankable project will increasingly be the one that can answer the following questions before the utility or RTO has to ask them.

Question	Why it matters
Can the load curtail in defined MW blocks?	Turns a binary reliability risk into a controllable operating resource.
Is dedicated or proximate supply identified?	Reduces resource-adequacy exposure and supports faster study treatment.
Who pays for upgrades?	Cost causation is becoming explicit; opaque rate-base subsidy is politically weaker.
Can service be staged?	Phased energization can bridge the gap between customer demand and grid construction.
How are BTMG and netting treated?	The economics of on-site generation depend on tariff treatment, not just engineering.
Are existing agreements protected?	Safe-harbor language will determine which projects avoid retroactive tariff disruption.

Practical screen: From this point forward, large-load projects will be underwritten on flexibility, cost responsibility and energy credibility as much as on tenant demand.

TOP STORIES

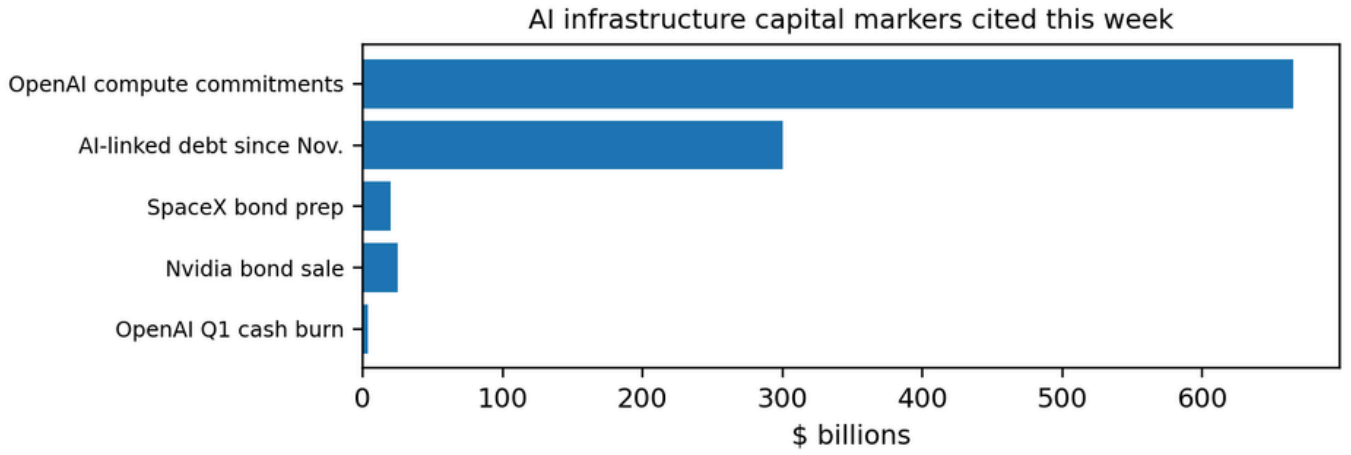

Links: Bloomberg / Nvidia bonds | Bloomberg / SpaceX bond sale | The Information / OpenAI cash burn | SemiAnalysis / capacity cancellation rebuttal | Bloomberg / FERC intervention

CAPITAL MARKETS ARE STILL FUNDING AI - BUT THE NUMBERS NOW LOOK LIKE INFRASTRUCTURE, NOT SOFTWARE

Nvidia's $25 billion high-grade bond sale, which reportedly drew as much as $85 billion of orders, is the cleanest signal that the credit market still wants AI exposure. The offering was upsized from an initial target of about $20 billion and followed a wave of AI-related borrowing from large technology companies funding data centers, semiconductors and related infrastructure.

SpaceX is the more aggressive version of the same theme. Bankers are preparing a bond sale of at least $20 billion to refinance a bridge loan after the company's record IPO, while analysts are modeling a capital program that could exceed $1 trillion by decade-end and potentially add more than $400 billion of net debt by 2031. That is not venture capital. It is railroad-scale infrastructure finance, with all the duration, execution and governance questions that implies.

OpenAI provides the operating-side caution. The company burned $3.7 billion in Q1 2026 against $5.7 billion of revenue, ended the quarter with more than $73 billion of cash and marketable securities, and disclosed an estimated $665 billion of compute spending commitments through 2030. Gross margins improved to 39%, but the core point remains: AI revenue is growing quickly, yet compute obligations are growing into a balance-sheet event.

Read-Through: The financing window is open, but it is becoming more selective. The market will fund platforms that look like critical infrastructure; it will also demand clearer proof that compute commitments can become durable cash flow rather than permanent capex absorption.

CAPITAL STACK READ-THROUGH

- $25B Nvidia bond sale; orders reportedly reached roughly $85B.
- $20B+ SpaceX bond expected to refinance bridge debt after the IPO.
- $300B+ AI-linked debt issuance since November across multiple credit markets.
- $665B OpenAI compute commitments through 2030; Q1 cash burn of $3.7B.
- Equity upside may absorb more risk than debt; credit investors will focus on duration, cash conversion and downgrade risk.

Links: Bloomberg / Nvidia $25B bond | Bloomberg / SpaceX $20B bond | Bloomberg / SpaceX AI fundraising | The Information / OpenAI Q1 burn

THE POLICY REGIME IS HARDENING: DATA CENTERS ARE NOW GRID, ENVIRONMENTAL AND NATIONAL-SECURITY POLICY

As outlined above, FERC's June 18 orders are the central policy event of the week. Grids covering roughly 65 million Americans and 21% of U.S. GDP must report within 30 days on generation adequacy for new and existing large loads, and within 60 days they must defend current tariff structures or propose changes. The public framing is speed; the commercial substance is cost causation, flexibility and resource adequacy.

The same tightening is happening in adjacent venues. CAISO is proposing computational-large-load ride-through rules that would require large compute loads to remain online through specified frequency and voltage events. Minnesota Power is defending Google load with a package that includes 10% demand response, 300 MW of wind and 400 MW of batteries. Virginia is preserving data-center incentives but proposing an electricity-use tax capped at $600 million per year. Together, these are not isolated regulatory stories. They are the emerging admission ticket for AI load.

The legal frontier is even more sensitive. The DOJ/xAI filing argues from national security and executive-branch control over environmental enforcement, while Commerce's order against Anthropic tests whether use of a frontier model by foreign persons can be treated as a controlled technology-transfer issue. Both cases are likely to be contested, but the message is clear: the AI stack is being pulled into national-security doctrine.

Read-Through: Regulatory support for AI infrastructure is not deregulation in the simple sense. It is a more interventionist bargain: faster approvals and strategic treatment in exchange for transparency, flexibility, cost responsibility and more political scrutiny.

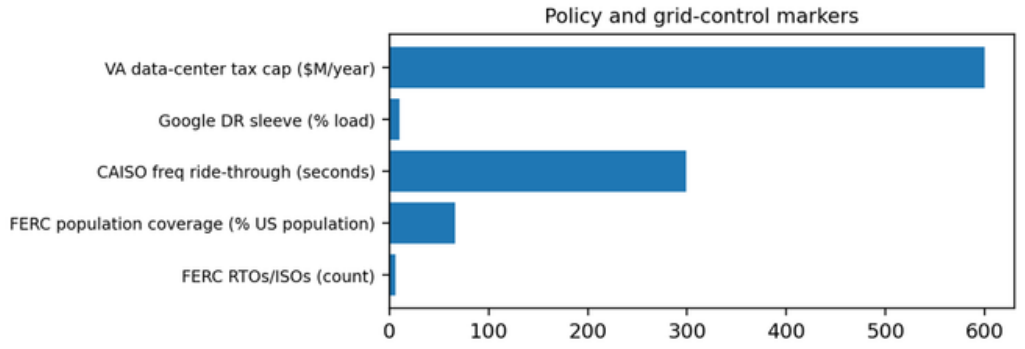

Links: Bloomberg / FERC intervention | **CAISO computational-large-load proposal** | **Minnesota Power-Google filing** | New York Times / DOJ-xAI case | **Bloomberg / Anthropic export-control order**

THE CANCELLATION NARRATIVE IS TOO BLUNT. THE REAL STORY IS PIPELINE TRIAGE.

"SemiAnalysis'" rebuttal to the "half of 2026 capacity is canceled" narrative is important because it sharpens the denominator. There are real delays - permitting, transformers, gas pipelines, interconnection queues and local opposition all matter. But treating every press-release megawatt as committed 2026 capacity creates a false cancellation story when the market later separates speculative announcements from executable projects.

The more useful framework is pipeline triage. Hyperscaler self-build forecasts and colocation forecasts have not moved dramatically in the "SemiAnalysis" model over the last six months, but individual speculative projects are slipping or being repriced. That squares with the rest of the week: Meta continues to sign compute deals; Tesla is exploring modular AI infrastructure concepts; legacy power sites are being revalued; and FERC is creating a process to distinguish serious large-load projects from generic load requests.

That distinction matters for valuations. The market should not pay the same multiple for announced GW, optioned land, permitted shells, utility-backed load, contracted capacity and energized power. Those are different risk categories. The week's development, regulatory and capital-markets news all points to the same underwriting rule: evidence beats narrative.

Read-Through:The AI buildout is not collapsing; it is being sorted. Execution-ready capacity should become more valuable as speculative capacity is discounted.

PIPELINE TRIAGE SCREEN

- **Demand evidence:** signed tenant contracts, not market rumors.
- **Power evidence:** utility agreement, resource plan, fuel strategy, interconnection status and curtailment rules.
- **Physical evidence:** permits, equipment slots, construction progress and visible critical-path work.
- **Financial evidence:** committed capital, risk-sharing deposits, offtake support and insurance capacity.
- **Political evidence:** local tax base, water/noise controls, ratepayer protection and community acceptance.

Links: SemiAnalysis / Stop Saying Half of 2026 US Datacenter Capacity Is Canceled | **Bloomberg / Crusoe-Meta** | **Bloomberg / FERC**

THE GEOPOLITICAL READ-THROUGH: AI ALONE IS NOT ENOUGH; POWER AND MANUFACTURING ARE THE COMPETITION

The China and industrial-policy pieces make the week's broader lesson more explicit. Bloomberg's essay argues that the US cannot win the technology race on AI alone; leadership also depends on drones, biotechnology, quantum computing, advanced manufacturing, supply chains and workforce diffusion. RAND's report reaches a complementary conclusion from the other side: China's Xi-era techno-industrial strategy is centralized, security-linked and finance-driven, with the Party-state mobilizing fiscal, financial, real-economy and overseas channels around priority industries.
For Digital Power, the implication is direct. AI infrastructure is not just a demand source for electricity. It is part of a national industrial-capacity race. China's nuclear buildout, oil-demand flexibility, industrial parks, standards strategy and manufacturing depth show what it means to treat energy and technology deployment as one system. The U.S. response does not need to copy China's model, but it does need to recognize that software

Read-Through: The strategic question is no longer "who has the best model?" It is "who can turn models into a scalable industrial system with power, hardware, manufacturing, financing and political legitimacy?"

INDUSTRIAL-COMPETITION MARKERS

- China leads in 57 of 64 Australian Strategic Policy Institute priority technologies, versus US leadership in 61 of 64 roughly two decades earlier.
- China manufactures roughly 70%-80% of global drones; the US has a stronger innovation base but thinner scale manufacturing.
- US biotechnology may generate more than $4T of global value over the coming decade, but scale-up manufacturing remains the constraint.
- Advanced manufacturing diffusion is the domestic challenge: moving even 20% of 250,000 SME manufacturers to the frontier means transforming 50,000 firms.

Links: Bloomberg / Why AI Alone Is Not Enough | RAND / China techno-industrial strategy | EIA / China nuclear buildout | WSJ / China oil shock absorber

DEVELOPMENT

THE DATA-CENTER LAND RUSH IS MOVING FROM "WHERE CAN WE BUILD?" TO "WHERE CAN WE CONTROL POWER?"

The development market is no longer simply aggregating acreage. It is aggregating entitlement optionality, power optionality, water narratives, local political consent and execution credibility. That is the read-through from this week's Salem Township, Wyoming, Maryland and Crusoe/Meta development news.

In Salem Township, the 4-3 Glen Brook Group announced a second proposed data-center land assembly that reads less like a real-estate transaction and more like an integrated infrastructure development thesis. The proposal covers more than 4,000 acres, contemplates a 39-building campus, includes a closed-loop water system, assumes on-site generation, and carries a potential land purchase price above $1.2 billion. Most importantly, the demand figure is being discussed at roughly 3.5 to 4.0 GW of potential power demand.

That number is the story. In the last cycle, a large land aggregation was valuable because it could be permitted, flattened, fibered and connected. In this cycle, the land is only valuable if it sits inside a close-by generation system that can plausibly support it. Salem's pitch is blunt: the township already hosts major transmission, substations, gas rights-of-way, Susquehanna nuclear infrastructure and the Moxie Freedom combined-cycle plant.

Wyoming is moving through the same filter. Uinta County commissioners unanimously approved Prometheus Hyperscale's 1.25 GW flagship project near Evanston, located on roughly 506 acres of ranchland. Prometheus is emphasizing that it will not draw cooling water from the nearby Bear River and will use a water-glycol cooling mixture changed out on a multi-year cycle.
Maryland offers the brownfield version of the same thesis. TeraWulf's proposed Chesapeake Data Campus would redevelop the former Morgantown coal plant with natural gas generation and a large battery, creating a data-center campus built around an existing energy asset rather than a greenfield utility request.

Links: Times Leader | KPCW | **The Baltimore Banner**

META'S 1.6GW CRUSOE DEAL SHOWS DEMAND IS STILL THERE — BUT COUNTERPARTY CONFIDENCE IS NOW PROJECT-SPECIFIC

Meta reportedly signed new agreements to buy AI compute capacity from Crusoe at two sites in Childress, Texas and Warrenton, Missouri, with combined capacity of roughly 1.6 GW. That is a major demand signal, particularly given Meta's broader commitment to spend at least $600 billion on AI infrastructure and its 5 GW Louisiana campus effort.

But the more interesting point is not that Meta needs more compute. Everyone knows that. The more interesting point is that hyperscalers are becoming far more selective about which megawatt pipeline they trust.

Crusoe is simultaneously proving both sides of the market. On one side, it continues to win hyperscale demand. Bloomberg reported that Crusoe now has 4.9 GW of contracted capacity and more than 40 GW in its total project pipeline. On the other side, it has faced reported execution friction, including being pushed out of the Wyoming project after Google raised cost and timetable concerns.

That juxtaposition is the development lesson: hyperscaler demand is not disappearing, but it is becoming more discriminating. Customers will sign for power-backed compute where they believe the developer can execute. They will walk, renegotiate or shift counterparties where the schedule, power plan or cost stack loses credibility.

Read-Through: The market should stop treating hyperscale demand and development execution as the same variable. The Meta/Crusoe news says demand for AI compute remains enormous. The Wyoming reporting says customers will not underwrite every megawatt equally.

Links: Bloomberg

THE NEW DEVELOPMENT STACK: RETROFIT COOLING, FLEXIBLE LOAD AND ON-SITE ENERGY ARE BECOMING PART OF THE REQUIRED STACK

Development strategy is moving inside the building. The old model was sequential: secure land, secure power, build shell, fit out racks. The new model is integrated: cooling architecture, demand response, battery strategy, on-site generation and flexible interconnection design are now part of the development thesis from day one.

Google's Brazos liquid-cooling sidecar is designed to bring liquid-cooled servers into legacy air-cooled data halls without full-building chilled-water retrofits. It can cool an adjacent 60 kW rack, deploy one rack at a time, use a closed-loop liquid-to-air system and avoid re-plumbing the building. That is not just an engineering update. It is a development unlock.
The same logic is showing up on the grid side. Minnesota Power's Google data-center filing says the utility is seeking approval for additional resources including 300 MW of wind, 400 MW of battery storage and up to 10% of Google demand available for demand response. This is the shape of the future large-load deal: not just serve the load, but serve the load with a defined flexibility envelope.

Verse is attacking the same bottleneck from the software side. The Nvidia-backed startup is partnering with Calibrant Energy to combine software, batteries and solar in ways that reduce peak grid burden and potentially accelerate data-center connection. The caveat is that utilities still lack a uniform framework for how batteries, on-site generation or flexible demand translate into faster service.

Links: Data Center Richness | Minnesota Power PUC response | Bloomberg | Construction Physics

ELECTRICITY MARKETS

THE LARGE-LOAD QUEUE IS BECOMING A GAS QUEUE, TOO

The power market is beginning to discover that data-center interconnection constraints do not disappear when the load moves behind the meter. They migrate. The bottleneck shifts from utility transformer queues and transmission studies to gas laterals, distribution-system reinforcement, firm fuel service, emissions permitting and construction lead times.

Enbridge Gas Utah and Wyoming has disclosed seventeen separate inquiries related to data-center electric generation or onsite electric generation for an unknown purpose. Three agreements have already been signed: two are in service and a third has an expected in-service date in the fourth quarter of 2026. Enbridge is explicit that it is unlikely all seventeen inquiries will be built, because some customers have walked away after seeing the estimated costs.

But the full-build number is still the signal: if all seventeen inquiries were constructed, estimated gas load would be approximately 4 Bcf per day, with typical project timelines of 18 to 24 months. In power-market terms, 4 Bcf per day is not a rounding error. At a rough 7.0 MMBtu/MWh heat rate, that fuel volume could support approximately 24 GW of continuous simple-cycle-equivalent electric output before derates, auxiliary load and operating constraints.

That does not mean the load will materialize. It means the new shadow queue is forming in the gas system. For data centers trying to bypass electric interconnection delays with onsite generation, the critical path becomes gas capacity, meter station design, pressure, rights-of-way, air permits, turbine delivery, water and local acceptance. Many developers may have left the electric queue, but they've entered a different queue with different constraints or they are pursuing both (the stronger strategy).
This matters because utilities, regulators and investors are likely to start underwriting large-load projects across both systems. A data-center campus that claims to be self-powered still has to prove its fuel. A BTM gas campus without firm supply, lateral economics and emissions clearance is not self-sufficient. It is just another speculative load with a different bottleneck.

ELECTRICITY MARKETS SNAPSHOT

- 17 inquiries — Enbridge Gas Utah/Wyoming has received data-center or onsite-generation gas service inquiries.
- 4 Bcf/d full-build case — Potential gas load if all seventeen inquiries were built, even though Enbridge says that outcome is unlikely.
- +40% RGGI auction move — Virginia re-entry and AI load growth helped drive a step-change in allowance pricing.
- $13/month rider — Dominion requested a customer charge to recover carbon-allowance costs.

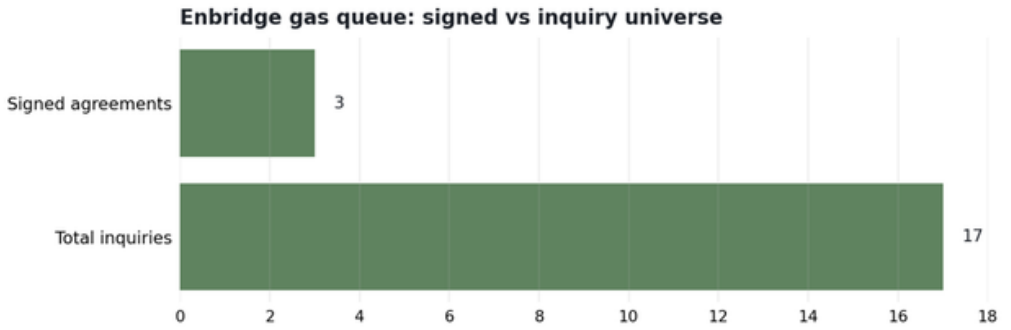

Source: Enbridge Gas Utah/Wyoming IRP excerpt. Three signed agreements out of seventeen inquiries; full-build case equals roughly 4 Bcf/d.

Links: Enbridge Gas Utah/Wyoming IRP excerpt

VIRGINIA'S DATA CENTERS ARE NOW A CARBON-MARKET INPUT

Data Center Alley is no longer just a PJM load story. It is becoming a carbon-market story. Virginia's return to the Regional Greenhouse Gas Initiative (RGGL) under Governor Abigail Spanberger comes as AI-driven demand growth is pushing traders to price tighter allowance supply, higher generator compliance costs and the possibility that those costs flow through to ratepayers.

Bloomberg reported that Virginia is rejoining RGGI after leaving in 2022, while the surge in power demand tied to data-center development has helped drive a rush into the market. In the most recent auction, allowance prices jumped 40% above the March sale, and the secondary market traded 21% above the auction price - an all-time record gap. That is the market saying the data-center demand curve has collided with a shrinking allowance supply curve.

The political consequence is immediate. Dominion Energy has asked regulators to approve an additional $13 monthly charge to reflect allowance costs. That number will become the public face of the issue. Data-center load growth may be a regional economic-development engine, but when it tightens carbon markets and shows up on residential bills, the politics change from growth to allocation: who pays, who benefits and who gets protected.

The supply side is tightening as well. RGGI has announced that it will reduce allowances sold over the next decade from 69.8 million short tons in 2027 to nearly 10 million by 2037. That cap trajectory creates a structural scarcity backdrop even before Virginia load growth is fully reflected. Offsetting forces exist - including new hydro imports from Quebec into New York through the Champlain Hudson Power Express and potential RGGI policy changes to contain prices - but traders are not yet convinced there is enough supply to absorb the new demand.

The larger read is that electricity markets are becoming multi-commodity. A data center is not just a load node. It can move capacity prices, transmission plans, gas-system investment, emissions allowances, utility riders and state politics. The market is starting to price that complexity.

Read-Through: The lesson is that AI load must be framed as a system-benefit asset, not merely a private offtake opportunity. Projects that bring controllable load, onsite generation, DR, local tax base and grid support will be easier to defend than projects that simply increase scarcity and push costs into regulated rates.

Links: Bloomberg

FINANCE

FINANCE SNAPSHOT

- **$54B —** US-listed companies have issued convertible bonds YTD, up 43% from the same period in 2025.
- **$73B+ —** Global hybrid-bond issuance in dollars, euros and sterling YTD - fastest pace in at least a decade.
- **$7.4B —** DeepSeek first funding round, valuing the Chinese AI lab at more than $50B.
- **$13B —** Baseten proposed valuation ceiling in a $1.5B dual-tier fundraising.
- **$134B —** Estimated cumulative global data-center insurance premiums from 2026-2030.

AI FINANCE IS STILL IN THE BOOM PHASE - BUT THE CAPITAL STACK IS GETTING MORE SOPHISTICATED

The AI trade is still being described by investors as a boom rather than full euphoria, but the financing architecture underneath it is becoming more complex. Bank of America's June fund-manager survey found that 56% of respondents called the current AI cycle a boom, while 21% called it euphoria and 9% described it as profit-taking. That is a market that knows it is crowded, but has not yet decided it is broken.

The warning label is concentration. Eighty percent of respondents said buying and holding global semiconductor stocks was the most crowded trade, the highest reading in the survey's history. At the same time, investors only trimmed their technology overweight to 26% from 33% and their global equity overweight to 38% from 50%. In plain English: investors are nervous about crowding, but not nervous enough to leave.

That matters because the financing channels are widening. The AI trade is no longer just venture equity and public-market multiple expansion. It now includes convertible bonds, hybrid debt, asset-backed financing, private credit, data-center securitization, insurance-linked risk transfer and utility balance-sheet engineering. The trade is still risk-on, but it is institutionalizing.

The market is therefore moving from a simple bubble question - are AI equities too expensive? - to a better capital-markets question: can the AI ecosystem keep finding cheaper, deeper and longer-duration capital to fund the physical infrastructure required to make the revenue story real? For now, the answer appears to be yes. But the cost of that answer is more leverage, more structural complexity and more dependence on continuously open capital markets.

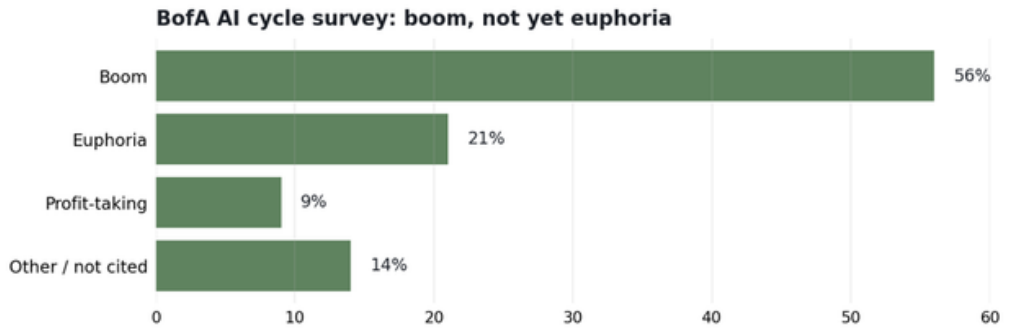

Source: Bank of America fund-manager survey reported by Bloomberg.

Links: Bloomberg / BofA survey

CONVERTIBLES AND HYBRIDS ARE BECOMING THE AI INFRASTRUCTURE FUNDING VALVES

The AI boom is pushing borrowers into securities that sit between clean equity and plain-vanilla debt. US-listed companies have issued about $54 billion of convertible bonds this year, up 43% from the same period in 2025 and the highest year-to-date level since the start of the Covid-19 pandemic. CoreWeave recently issued $4 billion of convertibles at a 1.75% coupon, while Akamai issued $3.5 billion of zero-coupon convertible senior notes.

The appeal is straightforward. AI issuers want growth capital without immediately issuing common equity. Convertible investors accept low or even zero coupons because volatility gives them equity optionality. That is the core bargain: the issuer monetizes its stock volatility; the investor receives downside bond structure plus upside participation if the AI trade keeps working.

Utilities are solving the same capital problem with a different instrument. NextEra has emerged as the largest hybrid-bond issuer globally after selling three hybrid bonds, including a structure with a 40-year maturity and a 20-year first-call date. Hybrid bonds are subordinated instruments that can receive partial equity credit from rating agencies, allowing utilities to fund data-center-driven capital needs without full common-equity dilution.

Hybrid issuance in dollars, euros and sterling has exceeded $73 billion year-to-date, nearly five times the comparable 2023 level. The driver is not hard to see: utilities face a step-change in grid, generation and large-load spending, while investors are looking for high-grade yield and diversification away from pure technology credit exposure.

The risk is that both structures depend on market tone. Convertibles need high stock prices, tight credit spreads and equity volatility. Hybrids need investor comfort with duration and subordinated risk. If AI sentiment turns, these funding valves can narrow quickly.

Read-Through: AI infrastructure capital is migrating into every corner of the capital stack. AI platforms should expect investors to compare project finance, converts, preferred equity, hybrid-style instruments and private credit against each other. The winning structure will be the one that funds power development without giving away too much upside or stressing the balance sheet.

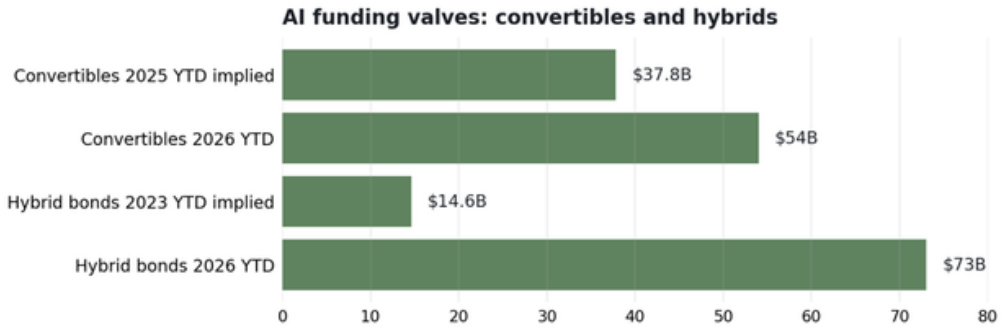

Source: WSJ/Dealogic and Bloomberg; implied prior-period values calculated from reported growth multiples.

Links: WSJ / convertibles | **Bloomberg / NextEra hybrids**

CSQUARE IPO TESTS WHETHER PUBLIC MARKETS WILL PAY FOR COLOCATION SCALE AGAIN

Brookfield-backed Csquare has filed for an IPO and plans to list on the NYSE under ticker CSQR. The filing is important because it tests whether the public market wants another scaled colocation infrastructure vehicle at a moment when AI power demand is expanding but investors are increasingly focused on leverage, asset quality and deliverability.

The operating base is real. Csquare owns and operates 64 data-center sites across 21 major metropolitan markets in the US, Canada and the UK, excluding three sites slated for closure. As of March 31, the platform had approximately 389 MW of sellable power capacity and more than 36,600 interconnection products. For the March quarter, revenue increased to $270.5 million from $232.8 million, adjusted EBITDA rose to $108.3 million from $86.3 million, and net loss widened to $66 million from $34.9 million.

The valuation marker is the key. Earlier reporting indicated that Csquare generated about $500 million of EBITDA and could be worth at least $8 billion, with some valuation expectations reaching $10 billion to $12 billion. That implies an EBITDA multiple range of roughly 16x to 24x, depending on where the deal prices and how investors underwrite growth, churn, capex and power scarcity.

This is not a pure AI factory IPO. It is a colocation and interconnection platform with AI optionality. That may actually be the point. Public investors may be more receptive to an infrastructure-like platform with existing cash flow, interconnection density and Brookfield sponsorship than to a pure-play development story with a longer path to stabilized EBITDA.

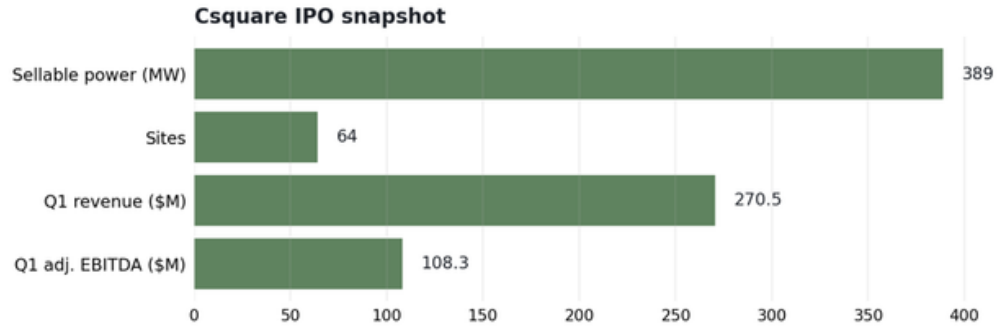

Source: The Deal; company filing figures as reported.

Links: The Deal / Csquare IPO

OPEN-SOURCE AI IS RAISING LIKE INFRASTRUCTURE - BECAUSE INFERENCE COST IS BECOMING THE PRODUCT

The AI funding cycle is broadening beyond the frontier-model labs. DeepSeek reportedly closed a funding round of more than 50 billion yuan, or about $7.4 billion, valuing the Chinese AI lab at more than $50 billion. Baseten is finalizing a $1.5 billion round at a dual-tier valuation of $11 billion to $13 billion. Both deals point to the same shift: open-source and inference infrastructure are now investable at hyperscale valuations.

DeepSeek's structure is unusually control-heavy. Investors are required to invest through a limited partnership managed by CEO Liang Wenfeng rather than directly into DeepSeek, and outside investors face a five-year lockup with no voting rights. Liang is reportedly contributing 20 billion yuan personally, Tencent 10 billion yuan and CATL 5 billion yuan, while China's national AI fund is the exception that invests directly and receives voting rights.

Baseten is the US-market counterpart to the same thesis. The company provides software and compute capacity to help enterprises run, optimize and train lower-cost open-source models. Its platform sources capacity from 20 different cloud providers. The business exists because the quality gap between open-source and closed-source models has narrowed, while enterprise customers are increasingly sensitive to escalating AI costs.

The strategic implication is that inference is becoming a cost-optimization market, not just a model-quality market. If enterprises can get good-enough results from open-source models with lower-cost infrastructure, value shifts toward orchestration, deployment, routing, optimization and compute access. That still requires power and data-center capacity - but it may change where the highest-margin layer sits

Read-Through: Investors should watch the inference ecosystem closely. If open-source inference scales, demand may become more distributed, cost-sensitive and latency-sensitive than frontier training demand. That will open up a market for modular, power-backed campuses that can serve inference loads without requiring in addition to the large single-campus training clusters.

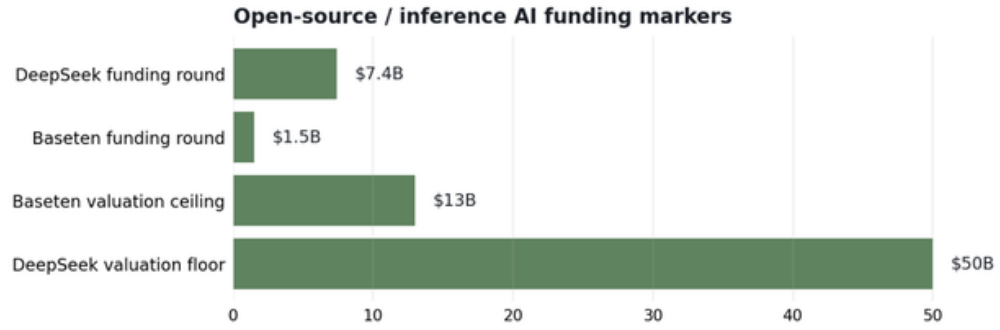

Source: The Information and WSJ; valuation and funding markers as reported.

Links: The Information / DeepSeek | WSJ / Baseten

DATA-CENTER RISK IS BECOMING AN INSURABLE ASSET CLASS

The data-center boom is now large enough to stress the insurance market. Zurich Insurance says the average value of a data-center project in its portfolio has increased from $150 million five years ago to $3 billion today. That is a 20x increase in average project exposure, before considering the concentration risk created by GPU-heavy equipment, weather exposure, construction delay and private-credit performance thresholds.

Zurich's message is that traditional insurance capacity may not be enough. The company expects discussions around securitization and insurance-linked securities to become more necessary as data-center projects scale. Artemis estimates cumulative global insurance premiums associated with data centers could reach $134 billion between 2026 and 2030.

This is another sign that AI infrastructure is becoming a full capital-markets asset class. First came venture equity and hyperscaler capex. Then came project finance, private credit, converts and hybrid debt. Now the risk-transfer market is arriving. Insurance-linked securities would allow investors to absorb specific property or catastrophe risks without underwriting whether the data center itself is profitable.

The risk-transfer story also intersects with development. Weather, grid fragility, cooling failure and surrounding-infrastructure outages all matter more as campuses get larger. A $30 billion total insurable value at a single location turns resilience from an engineering detail into a financing condition.

Read-Through: Insurance capacity should be treated as part of the development stack. Sites with lower climate exposure, brownfield industrial resilience, controllable power and clear emergency planning may price better not only with tenants and lenders, but also with insurers.

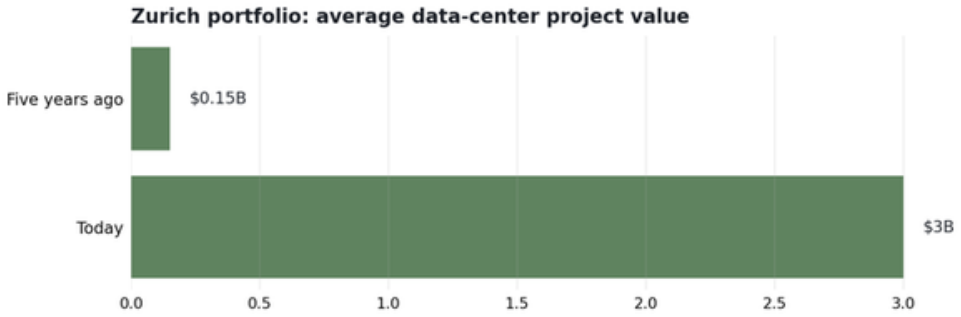

Source: Zurich comments reported by Bloomberg.

Links: Bloomberg / Zurich insurance securitization

INTERNATIONAL

CANADA'S SOVEREIGN AI STACK MOVES FROM POLICY SLOGAN TO PRODUCTION CONTRACT

BCE, Cohere, Hive Digital and Hypertec are turning Canada's sovereign-AI strategy into an operating stack. BCE will provide data-center capacity from its Merritt, British Columbia facility; Hive will deliver AI cloud services; Hypertec will provide Quebec-based hardware; and Cohere will use the compute to run models for government and enterprise customers. BCE CEO Mirko Bibic described the initiative as fully Canadian from top to bottom.

The dollar figures are not hyperscaler-sized, but they matter. Hive's contract is worth $220 million over three years, and BCE expects AI-powered solutions revenue to increase by at least C$500 million, or roughly $356 million, to C$2 billion by 2028. BCE is also building a 300 MW Saskatchewan data center in partnership with CoreWeave and Cerebras. Put together, the strategy is not just telecom diversification. It is a national compute-sovereignty platform built around domestic data centers, domestic hardware, domestic cloud services and a domestic model company.

The key market read-through is that sovereign AI is becoming a procurement and infrastructure strategy, not just an industrial-policy slogan. Countries that do not want to rely entirely on U.S. or Chinese compute will need domestic data centers, power, cooling, secure cloud, model access and procurement channels. That turns telecom companies, energy developers and local hardware suppliers into strategic AI infrastructure counterparties.

Read-Through: Sovereign AI can create a second buyer class for data-center power beyond hyperscalers: governments and regulated national champions that value domestic control, resiliency and data jurisdiction

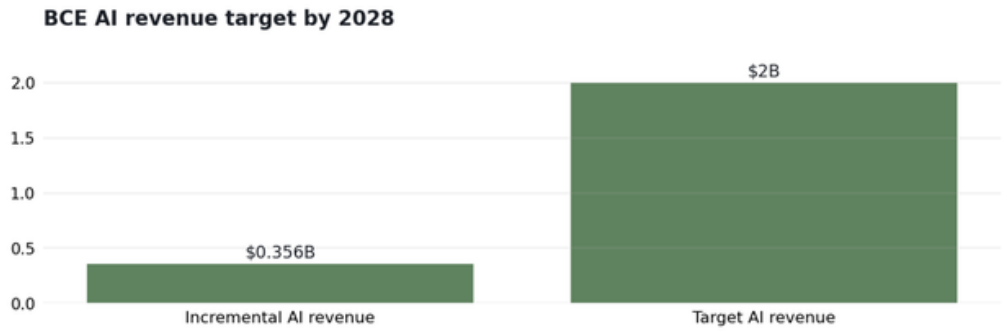

Source: Bloomberg; BCE target converted from reported C$500 million incremental revenue and C$2 billion 2028 AI-powered solutions revenue target.

Links: Bloomberg / BCE-Cohere

CHINA'S OIL IMPORT CUT IS HOLDING DOWN THE GLOBAL SHOCK - FOR NOW

The most important oil-market stabilizer during the Iran war may not be a new supply source. It may be China not buying. WSJ reports that China's crude imports fell to 7.8 million barrels per day in May, down from roughly 11 million barrels per day in recent years. That roughly three-million-barrel-per-day gap is about the combined daily oil consumption of Italy and France. It helps explain why Brent has stayed below $100 rather than moving toward the $150-$200 levels analysts feared under a prolonged Strait of Hormuz disruption.

The import cut appears to be coming from multiple levers. China had stockpiled cheap Russian and Iranian oil before the war, and users began drawing inventories at around 500,000 barrels per day in May. China also reduced transportation-fuel exports by roughly 500,000 barrels per day, cut refinery runs by 10 percentage points, and saw steam-cracker runs decline by 7 percentage points. On the demand side, oil is being displaced by electrified transport: May Day air passenger traffic fell 5.7% year-over-year, rail passenger traffic rose 4.6%, highway EV charging volume surged 53%, and about 15.4 million EVs traveled per day during the holiday period.

The risk is duration. Some of the decline is structural - EVs and rail are real oil substitutes. Some of it is temporary - reserve draws, refinery cuts and petrochemical curtailment can only last so long before feedstock shortages, producer-price inflation and export costs rise. If China has to return to the spot crude market at scale, the oil price shock could reaccelerate.

Read-Through: China's oil flexibility is an energy-security asset. The US should read this as a warning: electrification, fuel diversity and strategic reserves are not climate abstractions. They are macro shock absorbers.

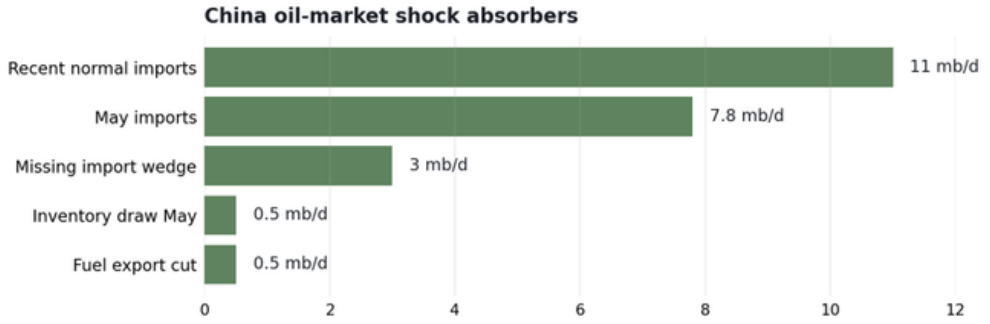

Source: WSJ, Vortexa, Argus and Chinese agency data as reported by WSJ. mb/d = million barrels per day.

Links: WSJ / China oil imports

CHINA'S NUCLEAR BUILDOUT IS WHAT INDUSTRIAL POLICY LOOKS LIKE WHEN IT REACHES THE POUR-CONCRETE STAGE

EIA's latest China nuclear update delineates that the power race is also an industrial-supply-chain race. China's operating nuclear capacity increased from 31.4 GW in 2016 to 58.7 GW as of May 2026. From 2016 through 2024, capacity increased 76%, or 24 GW, and China added another 1.1 GW in 2025 and 2.2 GW through May 2026.

The operating fleet now includes 60 reactors at 18 sites. More importantly, China has 36 reactors under construction across 19 sites that would add roughly 38.9 GW of additional capacity. EIA says that accounts for more than 49% of total world nuclear construction. The construction model is the point: standardized project management, batched reactor builds of six to ten units, domestic manufacturing of main plant components and build times of roughly five to seven years, compared with a global average closer to nine years.

This is the opposite of one-off nuclear development like Vogtle. China is treating nuclear as a repeatable industrial product. It is also moving into small modular reactors with Linglong-1, a domestically designed 100 MWe pressurized-water SMR intended for power generation, desalination and district heating. Whether or not the U.S. can replicate China's model, the implication for digital power is clear: long-term compute competitiveness will depend on countries that can build firm power at accelerated speed.

Read-Through: The U.S. AI buildout is running into power scarcity because demand has become industrial but supply remains procedural. China's nuclear data is a reminder that repeatable execution, domestic supply chains and standardized build programs are strategic weapons.

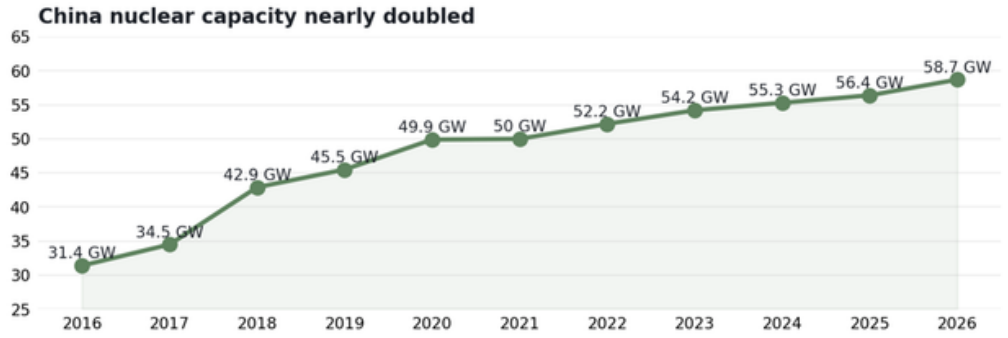

Source: U.S. Energy Information Administration and IAEA estimates; 2026 value through May.

Links: EIA / China nuclear capacity

OEMS

DETROIT'S EV HANGOVER IS TURNING INTO A GRID-STORAGE BUSINESS PLAN

The OEM story has flipped. Two years ago, the battery narrative was overwhelmingly about EV penetration, consumer adoption curves and IRA-supported gigafactories. Now the better incremental market may be the grid. GM is explicitly pitching itself as something closer to a distributed utility, while Ford is repurposing EV-era battery investments into grid-scale systems for utilities, data centers and industrial customers.

GM says more than 250,000 of its EVs on U.S. roads can already charge bidirectionally, meaning they can pull electricity from the grid and send it back. A firmware update is being rolled out for customers with GM Energy vehicle-to-home hardware, converting those systems into vehicle-to-grid assets without new hardware. The company is piloting the model with DTE Energy at 30 employee homes in Michigan and has sketched a 2030 PG&E scenario in which more than 52,000 GM EVs help balance the grid out of a projected 130,000 vehicles in the region.

The stationary-storage leg is just as important. GM and Peak Energy are partnering to develop and deploy sodium-ion battery systems for grid-scale storage, with GM developing sodium-ion cells in Michigan and receiving exclusive manufacturing rights. Peak says its sodium-ion systems could cut storage costs by 20% versus conventional systems and reduce energy waste from US grid batteries by up to 2 TWh per year under a wholesale switch from LFP. Peak already has 6.5 GWh of orders booked, including a Jupiter Power agreement that includes 720 MWh for 2027 and up to 4.75 GWh through 2030.

Ford is taking the more traditional OEM route: battery blocks rather than virtual power plants. Ford Energy is repurposing Michigan and Kentucky factories to produce LFP DC Block systems for utilities, data centers and industry, targeting at least 20 GWh of annual capacity and a five-year EDF North America framework for up to 20 GWh beginning in 2028. That is a direct response to the EV demand reset documented by the Dallas Fed: more than 20 US gigafactories announced in 2021-2022 represented over $50 billion of potential investment, but EV sales plateaued in 2025 after exceeding 1 million vehicles in 2023 and subsidy support faded.

Read-Through: The OEM read-through is that data-center load is becoming the second life of the EV supply chain. Battery manufacturing capacity that was overbuilt for consumer EV adoption can still become valuable if it is redirected toward substations, data centers, microgrids, demand response and backup power.

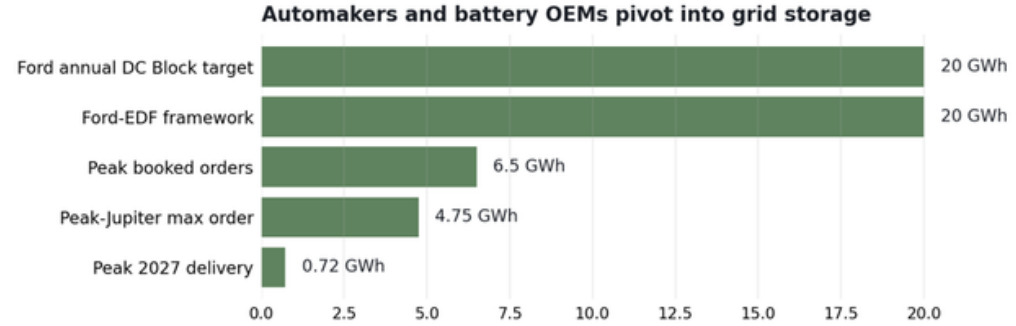

Source: Fortune, Utility Dive and Dallas Fed. Chart shows storage manufacturing targets, framework volume and order markers in GWh.

Links: Fortune / GM distributed utility | Utility Dive / Peak-GM | Dallas Fed / battery industry pivot

AI DATA CENTERS ARE TURNING SUPPLIERS INTO RISK MANAGERS, NOT JUST VENDORS

The AI supply chain is no longer taking hyperscaler demand at face value. Corning is the clearest example. The 175-year-old fiber and glass company has roughly doubled since January 1, is targeting a 50% sales increase by 2028, and has signed multibillion-dollar fiber deals with Nvidia, Meta and Amazon. Yet CEO Wendell Weeks is insisting on provisions that protect Corning if the world's largest tech companies do not ultimately need all the capacity they are ordering.

That is a major change in vendor posture. Corning's deal terms can include upfront customer funding to build the manufacturing capacity required to satisfy orders. Weeks' logic is simple: if the customer is creating the forecast risk, the customer should absorb part of the capacity risk. That lesson comes from the dot-com fiber bust, when suppliers were left holding the bag after customers overbuilt and demand collapsed.

Schneider Electric and Foxconn are moving in the same direction from the manufacturing side. Schneider said it will partner with Foxconn to develop AI data centers, with production beginning later this year. The collaboration is designed to help customers build and operate AI infrastructure with greater speed and efficiency - effectively moving the data-center supply chain closer to an industrialized, repeatable manufacturing model.

Together, Corning and Schneider/Foxconn show how the AI data-center supply chain is institutionalizing. Vendors want customer-funded capacity, risk-sharing, repeatable designs and manufacturing partners that can compress delivery schedules. The hyperscaler used to dictate terms. In constrained infrastructure markets, critical suppliers are beginning to demand terms of their own.

Read-Through: Supplier risk-sharing is a financing lesson. The more a project can convert tenant demand into deposits, take-or-pay commitments, equipment reservations and shared critical-path risk, the more credible the project becomes to lenders and equity investors.

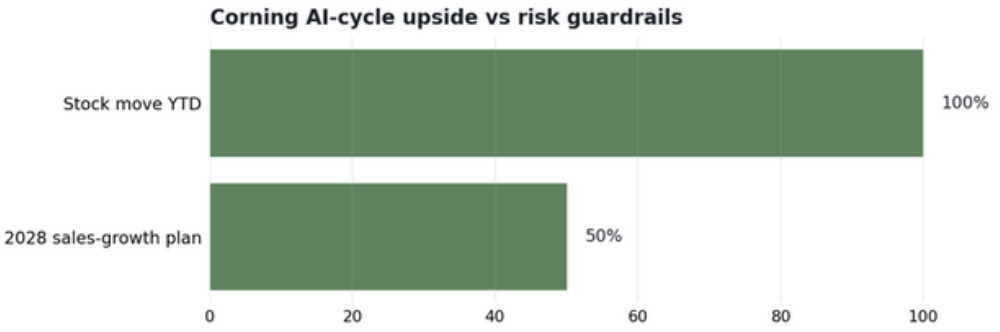

Source: WSJ. Stock move and sales-growth plan are approximate article figures.
Links: WSJ / Corning AI deals | WSJ / Schneider-Foxconn

ROLLS-ROYCE'S SWEDEN WIN PUSHES SMRS FROM POLICY ASPIRATION TOWARD EUROPEAN ORDERBOOK

Sweden selected Rolls-Royce to supply several small modular reactors for Vattenfall-backed new nuclear development, another tangible win for a European SMR vendor. Rolls-Royce beat GE Vernova in the final decision process, with Vattenfall citing the overall strength of the offer across supply chain, timing and returns. Rolls-Royce had already been selected by the UK and the Czech Republic, giving the company real momentum as European governments search for firm power options.

The Swedish project is meaningful in size. The plan calls for three reactors of 470 MW each, built partly in parallel, with construction on the second unit starting before the first is completed. The combined output is around 1,500 MW - slightly more than Sweden's largest reactor today - with a target to begin generating power by the middle of the next decade. The project company, Videberg Kraft, was set up by Vattenfall, with the state planning to own 60%, Vattenfall 20% and Industrikraft, a consortium of major power consumers, the remaining 20%.

The demand driver is industrial electrification. Sweden has six reactors in operation today, providing about one-third of national power, while hydro provides almost half. But the growth of weather-dependent renewables has increased price volatility, imbalances and grid-capacity constraints. The center-right government's answer is nuclear, even as critics argue storage and renewables can balance the system more cheaply.

The key OEM signal is that SMRs are being bought less as abstract clean-energy symbols and more as firm-power industrial tools. The buyer base is not only utilities; it includes governments and large power consumers that need long-duration certainty for electrification, industrial load and potentially compute.

Read-Through: *The SMR market remains timing-risked, but the Sweden award shows what bankability may look like: state support, utility sponsorship, industrial offtaker participation, repeatable module design and a defined mid-2030s energization target.*

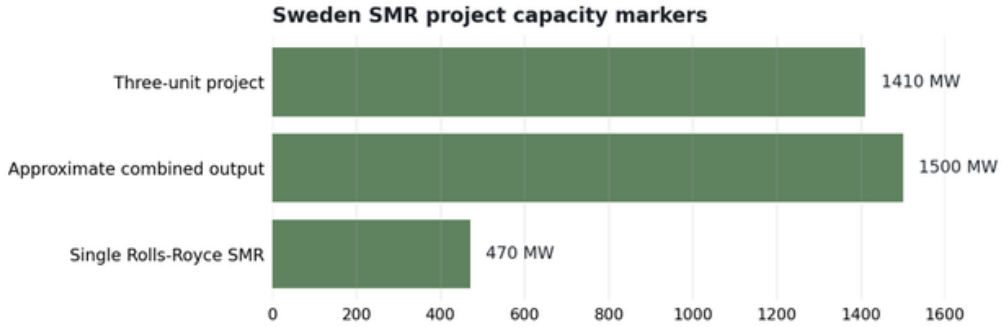

Source: Bloomberg. Capacity shown in MW; approximate combined output reported as around 1,500 MW.

Links: Bloomberg / Rolls-Royce Sweden SMRs

POLITICAL

VIRGINIA'S DATA CENTER COMPROMISE: KEEP THE INCENTIVE, TAX THE ELECTRICITY

Virginia's budget compromise turns the data-center backlash into a fiscal formula: preserve the incentive regime that built Data Center Alley, but tax the scarce input - electricity. The deal keeps the data-center tax incentives in place through their 2035 expiration while creating a new charge on electricity consumption rather than detonating the broader business-development compact.

The proposed fee is $0.011 per kWh of monthly electricity use, capped at $600 million annually, with collections above the cap refunded pro rata to operators. Virginia Scope reports $600 million of expected revenue in FY2027 and another $600 million in FY2028. Cardinal News frames the compromise against a larger backdrop: roughly $1.9 billion of annual foregone tax revenue tied to incentives versus roughly $9.1 billion of GDP support attributed to the sector.

The regulatory layer is just as important. The conference report directs Virginia DEQ toward data-center noise rules and cooling-water scarcity analysis, orders the State Corporation Commission to collect data on electric-service agreements, water usage and generator permitting, and creates a Joint Subcommittee on Tax Policy to keep studying the industry. The tax fight may be settling; the operating-regulation fight is beginning.

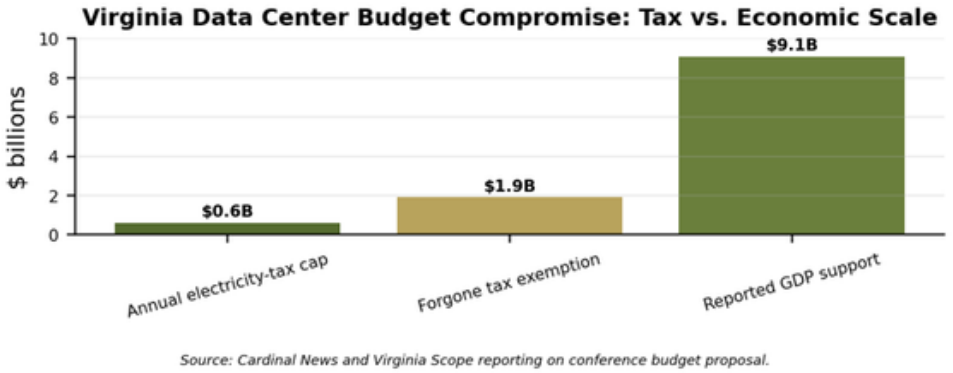

Source: Cardinal News and Virginia Scope. Figures reflect reported budget conference proposal, annual cap and GDP/tax-exemption context.

Links: Cardinal News / Virginia budget deal | Virginia Scope / $600M data center electricity tax

LOCAL ZONING IS BECOMING AI INDUSTRIAL POLICY

The Wall Street Journal's Loudoun-versus-Prince William comparison makes the local-politics point bluntly: AI infrastructure is being decided in county boards, courthouse appeals and zoning maps as much as in hyperscaler boardrooms. Loudoun created certainty early with a 2000 by-right zoning decision for data centers in commercial areas; AWS followed with its first county facilities in 2006.

The fiscal numbers explain why Loudoun remains the benchmark. Data centers are expected to generate nearly $1.3 billion of county taxes in 2027, or roughly 45% of the total, and the county has cut its residential property-tax rate by roughly 30% over the past decade. The political bargain is simple: high tax base, limited school burden, and infrastructure impacts that can be managed if planned early.

Prince William is the cautionary foil. A major project approved in 2023 effectively collapsed after a Virginia appeals ruling halted the project and local officials declined to defend it. Developers including Compass Datacenters withdrew after spending tens of millions. The broader lesson is that zoning instability can destroy even well-capitalized infrastructure plans - along with jobs, tax base and grid-investment leverage.

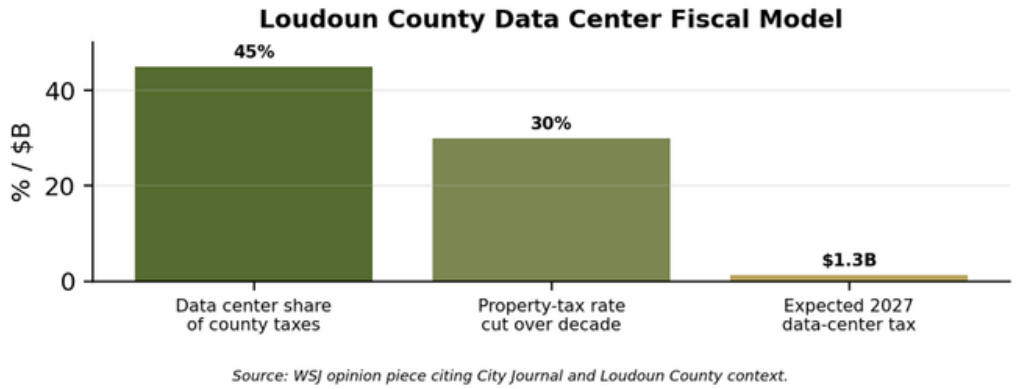

Source: WSJ opinion article citing Loudoun County fiscal figures and City Journal context; percentages and 2027 tax estimate are article figures.
Links: WSJ / Two Virginia Counties Diverge on AI Data Centers

REGULATORY

GRID CODES ARE COMING FOR DATA CENTERS: CAISO TURNS COMPUTE LOAD INTO A RELIABILITY ASSET

CAISO's straw proposal for "computational large loads" is a quiet but important regulatory marker: data centers are no longer being treated as passive retail customers. They are becoming grid assets with mandatory behavior during system disturbances.

The proposal's first-order message is ride-through. If system frequency rises above 61.2 Hz but stays at or below 61.8 Hz, or drops below 58.8 Hz but stays at or above 57.0 Hz, covered loads would need to continue operating for 299 seconds before tripping. Between 58.8 Hz and 61.2 Hz, the expected behavior is continuous operation. At the extremes — above 61.8 Hz or below 57.0 Hz — the load may trip.

Voltage treatment is similarly explicit. Between 0.9 and 1.1 per unit, continuous operation is required. If voltage falls below 0.9 per unit, the load must ride through for 6.0 seconds; below 0.7 per unit, 3.0 seconds; below 0.5 per unit, 1.2 seconds; and below 0.25 per unit, 0.15 seconds. That is not a commercial preference. It is a technical interconnection doctrine.

The practical effect is that large compute campuses will increasingly need inverter controls, staged load shedding, UPS/BESS coordination and facility-control software designed around the grid's disturbance envelope. "Fast trip" behavior may become unacceptable at scale because simultaneous data-center dropouts can worsen the same reliability event they are trying to survive.

Read-Through: *CAISO is a preview of where PJM, ERCOT and utility tariffs are heading. The most valuable AI-load projects will be those that can prove they ride through disturbances, shed load in controlled blocks, provide demand response and support grid stability rather than creating a single large binary trip risk.*

Links: CAISO computational-large-load excerpt

CAISO RIDE-THROUGH SNAPSHOT

- Frequency 58.8-61.2 Hz: continuous operation expected.
- Frequency 61.2-61.8 Hz or 58.8-57.0 Hz: 299-second minimum ride-through before tripping.
- Voltage 0.9-1.1 per unit: continuous operation expected.
- Voltage below 0.9 per unit: 6.0-second ride-through; below 0.7 per unit: 3.0 seconds; below 0.5 per unit: 1.2 seconds; below 0.25 per unit: 0.15 seconds.
- Strategic signal: compute load is being pulled into grid-code compliance, not merely utility service contracting.

DOJ-XAI FILING TURNS DATA-CENTER PERMITTING INTO NATIONAL-SECURITY DOCTRINE

The Justice Department's intervention in the NAACP Clean Air Act lawsuit against xAI is one of the more aggressive regulatory developments in the AI-power buildout. DOJ told a federal court in Mississippi that xAI should be allowed to operate dozens of gas-burning turbines despite the lack of permits, arguing that the lawsuit threatens national security by seeking to interrupt power supply for AI innovation supporting Department of War military operations.

The factual stakes are significant. The NAACP suit challenges xAI and subsidiary MZX Tech's use of portable natural-gas turbines to power Grok-related data-center operations near the Tennessee-Mississippi border. The complaint alleges that xAI operates 57 gas turbines at the Colossus 2 data center without Clean Air Act pollution controls and seeks penalties of roughly $124,000 per day per violation plus an injunction requiring the turbines to stop operating.

The legal stakes may be larger than the emissions case itself. DOJ argued that the executive branch should have authority to stop environmental citizen suits over federal objection. That would strike at a long-standing enforcement mechanism embedded in the Clean Air Act. The tension is sharper because EPA had previously clarified that even temporary turbines are subject to permitting and pollution-control obligations.

This is the collision the AI buildout has been heading toward: emergency speed, national-security framing, temporary power, local pollution burdens and legacy environmental statutes all landing in one courtroom. If the national-security argument succeeds even partially, developers serving defense-linked AI workloads may seek broader emergency-power latitude. If it fails, temporary generation for large compute loads becomes a higher-risk permitting path.

Read-Through: The lesson is not that permitting can be ignored. It is the opposite: temporary generation, emissions controls, community impact and legal defensibility must be engineered upfront. Sites with existing permits, brownfield power infrastructure and a credible environmental record should screen better than emergency turbine deployments that invite litigation.

Links: New York Times / DOJ seeks to halt pollution lawsuit against xAI data center

XAI / DOJ REGULATORY STAKES

- 57 gas turbines alleged by NAACP at Colossus 2, near the Tennessee-Mississippi border.
- $124,000 per day per violation sought in penalties, plus an injunction against continued turbine operation.
- DOJ argument: national security and executive-branch control over Clean Air Act enforcement should override the citizen suit.
- Counter-risk: conventional Clean Air Act reading may favor NAACP, with appeal risk likely through the Fifth Circuit and potentially the Supreme Court.
- Developer lesson: temporary power is no longer a back-office workaround; it is a front-page regulatory exposure.

MINNESOTA POWER-GOOGLE SHOWS THE NEW LARGE-LOAD COMPACT: RENEWABLES, BATTERIES AND DR AS THE ADMISSION TICKET

Minnesota Power's response in its integrated resource plan dispute shows how utilities are beginning to defend hyperscale load additions. Clean Energy Organizations asked the utility to confirm that it was not seeking approval of new resources solely to serve Google's Hermantown data-center project. The response reframes the load as part of a broader, cleaner and more flexible resource package.

The key data point is flexibility: up to 10% of Google's demand is available for demand response. Minnesota Power is also seeking approval for 300 MW of wind generation and 400 MW of battery storage, while saying it is not requesting new natural-gas generation beyond the IRP Base Plan to serve existing customers. The Google electric service agreement is docket 26-159, tied to Harmony Group LLC and Project Loon.

This is the emerging regulatory bargain for AI load: large customers can still come, but the package increasingly needs renewables, storage, demand response, transparent resource planning and a credible answer to "who pays?" Utilities that can present the load as flexible and investment-supporting will fare better than those that present it as a sudden bilateral burden.

Read-Through: A data-center offtake story should be paired with a flexibility story: dispatchable curtailment, BTM storage, backup coordination or interruptible blocks are now part of the value proposition.

Links: Minnesota Power / Google demand-response note | Minnesota Power IRP docket 26-159 reference

APPENDIX A | MARKET DATA

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

Metric	Latest / This Week	Prior / Reference	DPW Read
Nvidia bond sale	$25B high-grade offering; ~$85B order book	First bond sale since 2021; upsized from ~$20B	AI supply-chain leaders are now funding ecosystem expansion through large public credit deals.
SpaceX bond plan	At least $20B contemplated first high-grade dollar bond	Refinances temporary $20B bridge loan due Sept. 2027; $29.1B long-term debt at Mar. 31	IPO was not the end of the funding cycle; it opened a debt-market phase.
SpaceX capex ambition	> $1T modeled through decade; possible >$700B/year capex by 2031 in one estimate	AI operations and orbital data-center vision drive estimates	The market is being asked to finance infrastructure before proof of normalized returns.
OpenAI Q1 cash burn	$3.7B cash burn on $5.7B revenue	39% gross margin; $73B cash/securities after March raise	Demand is strong, but model economics remain deeply compute intensive.
OpenAI compute commitments	$665B estimated spending commitments through 2030	Largely off balance sheet per shareholder materials	Long-duration compute obligations are becoming the real AI liability stack.
SemiAnalysis 2026 capacity check	YE2026 NA hyperscaler self-build forecast moved ~1%; NA colocation <5% over six months	Counters “half of 2026 capacity canceled” narrative	Delays are real, but the headline bears appear to be overstating executable-capacity risk.
FERC large-load orders	6 jurisdictional RTOs/ISOs under Section 206 show-cause pressure	PJM, MISO, SPP, CAISO, ISO-NE, NYISO	Federal policy is moving toward speed-for-responsibility: faster interconnection, more cost causation and flexibility.
FERC public target frame	90-day processing paradigm discussed publicly	Not yet a binding generic approval clock	Treat as directional policy pressure, not final tariff law.
CAISO computational loads	299-sec frequency ride-through; 6.0-sec voltage ride-through at <0.9 p.u. band	Straw proposal for computational large loads	Data centers are being regulated as reliability actors, not passive retail customers.
Minnesota Power / Google	Up to 10% of Google demand available for DR; 300 MW wind; 400	Project Loon ESA / IRP response	Flexible load plus new resources is becoming the admission ticket.
Virginia data-center tax	$0.011/kWh proposed electricity fee; $600M annual cap	Preserves sales/use-tax incentive through 2035	Politics is shifting from whether to host data centers to how to tax scarce inputs.
Loudoun fiscal model	~$1.3B expected 2027 data-center tax revenue; ~45% of county total	Data centers fund core services without proportional resident burden	Local fiscal design is becoming AI industrial policy.
Global DC energy / capacity	~415 TWh 2025 energy use; ~100 GW installed capacity; nearly +100 GW by 2030	JLL research cited in Development section	Global capacity may double, but power, land and permitting remain the gating variables.
Meta Hyperion	$200B+ spend; ~5 GW; up to 7,500 workers; 10 gas turbines	Richland Parish case study	Community-benefit architecture is becoming as important as interconnection architecture.
Rolls-Royce Sweden SMR	3 x 470 MW reactors; ~1.4-1.5 GW total program	Vattenfall-backed Swedish selection	SMRs are moving from policy narrative toward industrial order book.

APPENDIX B | RTO / FERC CALENDAR

Expanded calendar combines the prior DPW calendar with the June 18 FERC large-load Section 206 order package. Weekend/holiday due dates should be checked against FERC filing rules and any order-specific language.

Date	Event / Filing Clock	RTO / Agency	Why It Matters
Jun. 18, 2026	FERC large-load Section 206 show-cause orders issued to six jurisdictional RTOs/ISOs	FERC; PJM, MISO, SPP, CAISO, ISO-NE, NYISO	Start of the federal defend-or-revise process for large-load tariffs, cost causation, co-location, BTM generation and flexible service.
Jun. 18, 2026	FERC public framing: accelerate hyperscaler connection requests toward a 90-day processing paradigm	FERC	Important policy signal, but should not be treated as a universal binding approval clock until final tariffs are filed and accepted.
Jun. 18, 2026	PJM E-2 co-location / BTMG order: FERC accepted in part and rejected in part PJM compliance filing	FERC / PJM	Keeps PJM co-location, firm/non-firm contract-demand service and interim transmission service on a 60-day compliance path.
Jun. 18-Jul. 18, 2026	30-day generation-adequacy report clock	FERC / six RTOs/ISOs	Each market must explain how adequate generation will be available to serve existing and new large loads. July 18 is the calendar date; if weekend filing rules apply, confirm whether the effective date rolls to Jul. 20.
Jul. 17, 2026	Texas PUC/ERCOT MOU due under Abbott directive	Texas PUC / ERCOT	Keeps pressure on data centers and large loads to fund incremental infrastructure and avoid ratepayer cost shifting.
Jul. 18 / Jul. 20, 2026	30-day RTO/ISO reports due on generation adequacy for large loads	FERC / PJM, MISO, SPP, CAISO, ISO-NE, NYISO	First concrete deliverable from the June 18 order package; will reveal how each region frames resource adequacy and large-load risk.
Jul. 23, 2026	FERC AD26-7-000 technical conference / reform discussion	FERC / PJM / RTO stakeholders	Important venue for reliability backstop, resource adequacy, cost allocation and large-load policy framing.
Jul. 31, 2026	Texas residential transmission cost-cut deadline	Texas PUC / ERCOT	Ratepayer-protection marker as load growth and transmission cost allocation remain politically sensitive.

Aug. 17, 2026	60-day defend-or-revise tariff responses due from each RTO/ISO	FERC / six RTOs/ISOs	Core deadline: each market must justify current tariff treatment as just and reasonable or propose revisions for large-load service, study process and cost allocation.
Aug. 17, 2026	PJM further compliance filing due in E-2 co-location / BTMG proceeding	PJM / FERC	Key watch item for interim NITS, firm contract demand, non-firm contract demand, BTMG netting, minimum charges and curtailment priority.
Sep. 16, 2026	90-day policy milestone from June 18 order date	FERC / RTOs / utilities	Not a generic legal deadline, but a public benchmark for whether the federal large-load speed agenda is translating into operational changes.
Oct. 2026	PJM expected to identify first 10 EIT projects	PJM	Two months before the 2029/30 BRA, EIT selections will show which projects PJM views as near-term reliability-relevant.
Dec. 2026	PJM 2029/30 Base Residual Auction	PJM	Next major capacity-price signal for data-center-driven demand, retirements, storage, gas and demand-response economics.
Jan. 2027	Expected PJM Connect & Manage timing / cutoff watch item	PJM	Large-load interconnection treatment and curtailment exposure could shift materially depending on final PJM implementation.
YE 2027	PJM EIT track sunset under current temporary framework	PJM	Three-year COD window closes unless extended or replaced by a durable tariff architecture.
2028	Minnesota Power / Google resource package watch window	Minnesota PUC / Minnesota Power	Track IRP treatment of 300 MW wind, 400 MW battery storage and 10% Google demand-response flexibility.
2028-2030	Meta / Zelestra U.S. solar portfolio expected online	Meta / Zelestra / ISOs	Renewable procurement will be watched as political proof that new AI load can be paired with new generation.
2030	JLL nearly +100 GW global DC capacity buildout horizon	Global DC market	Long-run benchmark for land, power, equipment and grid-delivery pressure.

APPENDIX C | KEY LINKS

Resource	Live link
FERC Large-Load Integration Release	FERC - Large Load Integration
FERC Fact Sheet	FERC - Supercharge America's Grid
FERC June Commission Meeting Summaries	FERC - June 2026 summaries
SPP High Impact Large Load Integration	SPP HILL Integration
Bloomberg - FERC Interventionist Posture	Bloomberg - US power regulator pivot
Bloomberg - Nvidia Bond Sale	Bloomberg - Nvidia $25B bonds
Bloomberg - SpaceX Bond Sale	Bloomberg - SpaceX $20B bond sale
Bloomberg - SpaceX AI / Data Center Growth	Bloomberg - SpaceX fundraising campaign
The Information - OpenAI Cash Burn	The Information - OpenAI Q1 burn
SemiAnalysis - Data Center Capacity Rebuttal	SemiAnalysis - Stop saying half canceled
RAND - China Techno-Industrial Strategy	RAND - Producing Under Pressure
Bloomberg - AI Alone Is Not Enough	Bloomberg - US / China tech race
New York Times - DOJ / xAI Clean Air Act Case	NYT - DOJ seeks to halt xAI suit
Basenor - Tesla MEGAPOD Trademark	Basenor - MEGAPOD filing
Energy News Beat - Tesla MEGAPOD	Energy News Beat - Tesla AI infrastructure
Wall Street Journal - Virginia Counties Diverge	WSJ - Loudoun vs. Prince William

APPENDIX

APPENDIX D | GLOSSARY + WATCH LIST

Term	Definition
Section 206	Federal Power Act authority allowing FERC to find an existing rate, term or practice unjust and unreasonable and require a replacement.
Defend-or-revise	The posture of the June 18 large-load orders: RTOs/ISOs must justify existing tariffs or propose tariff revisions.
HILL	SPP High Impact Large Load process - a defined framework for studying and integrating very large commercial/industrial loads.
HILLGA	SPP High Impact Large Load Generation Assessment - a paired-generation pathway for resources serving high-impact load.
LLRIS	Load Limited Resource Interconnection Service - service concept limiting a generator's interconnection to the amount needed for
BYOP / BYONG	Bring Your Own Power / Bring Your Own New Generation - large-load customer arrives with dedicated or proximate supply rather than only
BTMG	Behind-the-meter generation, often central to co-location debates and netting/minimum-charge issues.
Contract Demand Service	PJM service concept for firm or non-firm transmission service up to a specified MW contract demand for co-located load.
Demand Response sleeve	Contractual portion of large-load demand that can curtail during stressed conditions; increasingly part of utility approval narratives.
Ride-through	Obligation for large load to remain online for defined frequency or voltage disturbances instead of tripping immediately and worsening system instability.

APPENDIX

APPENDIX D | GLOSSARY + WATCH LIST

Watch Item	Signal	Next Check
FERC 30-day reports	How each RTO/ISO defines generation adequacy for large loads	July 18/20, 2026
FERC 60-day tariff responses	Whether regions propose flexible service, cost-causation and co-location pathways	Aug. 17, 2026
PJM E-2 compliance	BTMG netting, minimum charges, contract demand, curtailment priority	Aug. 17, 2026
SPP HILL benchmark	Whether other regions copy HILL/HILLGA/LLRIS concepts	Summer/Fall 2026 filings
CAISO computational load rules	Whether ride-through becomes a broader North American template	CAISO straw proposal process
SemiAnalysis capacity revisions	Whether actual executable 2026 capacity changes materially rather than announcement pipeline only	Next model revision / earnings season
OpenAI / Anthropic IPO path	Whether public equity can finance compute obligations without margin shock	SEC / investor updates
AI credit issuance	Whether spreads stay tight as issuance grows	Nvidia, SpaceX, Oracle, CoreWeave, Cipher, Applied Digital
Virginia tax model	Whether electricity-use fees become a national model	Budget enactment / 2027 implementation
Local consent economics	Whether school funding, rate credits and local tax receipts become repeatable siting tools	Meta Hyperion / Loudoun / Texas and PA siting debates

APPENDIX

APPENDIX E | ARTICLE ROLLUP / SOURCE LIST

Section	Headline / Theme	Primary Source(s)	Date
Week in Review	AI infrastructure becomes industrial finance, policy and power problem	Bloomberg; RAND; SemiAnalysis; The Information	Jun. 15-20
Top Stories	Nvidia, SpaceX and OpenAI show the capital intensity of AI buildout	Bloomberg; The Information	Jun. 15-18
Top Stories	SemiAnalysis rebuts “half of 2026 capacity canceled” narrative	SemiAnalysis	Jun. 18
Top Stories	RAND: China’s techno-industrial strategy and U.S. industrial-base challenge	RAND	2026
Market Spotlight	FERC large-load Section 206 intervention; SPP reference, PJM test case	FERC; Bloomberg; SPP	Jun. 18
Development	Power-first AI campus design, brownfield sites and executable pipeline vs. announced GW	Bloomberg; SemiAnalysis; Google/Intersect; Meta; JLL	Jun. 2026
Electricity Markets	Storage gap, solar growth, fossil-retirement delay, energy-security premium	EIA/Utility Dive; Carlyle; The Register	Jun. 2026
Finance	AI infrastructure moves through IG bonds, bank loans, HY, private credit and IPO pipeline	Bloomberg; Apollo; WSJ; The Information	Jun. 2026
International	Canada sovereign AI, China power/nuclear and techno-industrial policy	Bloomberg; RAND	Jun. 2026
OEMs	Batteries, engines, SMRs, fiber and AI data-center supply chain risk sharing	Bloomberg; WSJ; Utility Dive; OEM releases	Jun. 2026
Political	Virginia tax compromise and Loudoun/Prince William local-policy split	Virginia Scope; Cardinal News; WSJ	Jun. 2026
Regulatory	CAISO ride-through, DOJ/xAI Clean Air Act, Minnesota Power/Google DR package	CAISO; NYT; Minnesota Power	Jun. 2026