KEEL INFRASTRUCTURE

DIGITAL POWER WEEKLY

ISSUE NO. 10
July 6 - July 12, 2026

EDITOR
Chris Ruppel
SVP - Power

~15 GW/YR
U.S. DATA CENTER GAP

10+ GW
ANTHROPIC CAPACITY BY 2026

7.5 GW
META HYPERION GAS FLEET

$71.9B
AI INFRA 144A SINCE MAY 2025

In this Issue

WEEK IN REVIEW

The Firm Megawatt Flywheel

There are weeks when the AI infrastructure story advances, and some when it clarifies. This week was the latter. On Thursday, the Catalyst podcast put Shayle Kann opposite SemiAnalysis's Jeremy Eliahou Ontiveros, and between them they analyzed the core Digital Power thesis that for an AI lab, power is revenue. Not an input cost, not a line in the site-selection matrix, but the binding constraint on the business itself. A frontier lab without megawatts is a research group with a waiting list.

Their arithmetic makes the point. Anthropic ended 2025 with roughly 1.5GW of capacity and intends to hold more than 10GW by 2027 — the equivalent of adding Google's entire global fleet in twenty-four months. Roughly half the gigawatts Amazon and Microsoft energize each quarter are now proxies for OpenAI and Anthropic; the labs, not the clouds, set the load curve and it significantly compounds each year against a grid expected to deliver against 50GW of demand only 35 GW of firm-equivalent generation in 2027–28 on the most generous math offered. No board commits fifty billion dollars per gigawatt of AI data center capex against utility timelines that carry no penalty for slippage. Behind-the-meter generation exists because certainty, not electricity, is the scarce commodity.

Meta then supplied the case study for the SemiAnalysis report out this week, one year into the Superintelligence Labs rebuild. The conclusion: models remain unremarkable but the inputs are anything but—five Titan clusters under simultaneous construction, capex guided to $125–145 billion, and three power procurement models run at once — pure off-grid gas at Prometheus, a dedicated 7.5 GW utility fleet in development at Hyperion, on-site bridge generation at El Paso.

The capital markets completed the flywheel circuit. Since May 2025, AI infrastructure borrowers have raised $71.9 billion through the 144A rail; the Hyperion/Blue Owl vehicle priced $27 billion of A+ paper at roughly 225 basis points over Treasuries; TeraWulf is pursuing $3.5 billion against a twenty-year, ~$19 billion Anthropic lease. The market is not declaring all AI infrastructure to be core — it is saying that investment-grade, take-or-pay, power-backed cash flow can now be capitalized like infrastructure, and the projects that finance first will outbid everyone else for turbines, land and talent.

Two caveats temper the enthusiasm. Wood Mackenzie sees Henry Hub migrating toward $5/MMBtu real by 2035 as AI load and LNG exports absorb the surplus — the molecule that solves the 2027 power problem could become a 2035 margin problem. And the social-license account is being drawn down: a Senate probe into gas-fired AI data centers, campaign-season defections on siting issues, and water disputes in Cheyenne, WY. BTM solves the queue; it does not solve legitimacy. The flywheel — megawatts become compute, compute becomes tokens, tokens become revenue, revenue funds the next megawatt — turns only where power is firm, financeable, and locally welcome.

BOTTOM LINE

Power is no longer an input to the AI business model — it is the rate limiter on it. The labs set the load curve, the grid cannot clear it, and the capital markets have now confirmed that contracted, power-backed cash flow prices like infrastructure. The projects that reach ‘financeability’ first will corner turbines, queue positions, and talent; everyone else will pay an increasing scarcity premium. But gas basis risk and eroding social license are structural threats, not footnotes. The scarce asset is firm, financeable, locally legitimate power on lab timelines — and whoever manufactures that certainty sets the terms for everyone else.

Catalyst Podcast Link: https://www.latitudemedia.com/news/catalyst-inside-the-ai-power-wars/

★ MARKET SPOTLIGHT | FEATURED ANALYSES

POWER IS REVENUE

The Catalyst interview frames the firm-megawatt flywheel

Kann's opening provocation was aimed at the digital-power investor class when he put forward that everyone can recite the silicon landscape, yet almost no one can articulate, company by company, the power strategies underneath the buildout. His second observation was the more striking: the war for AI research talent now has a quieter twin on the energy side, where the small population of people who genuinely understand digital power—interconnection, wholesale markets and on-site generation—has become a hot commodity and the movement of that talent is already redrawing the hyperscaler league table.

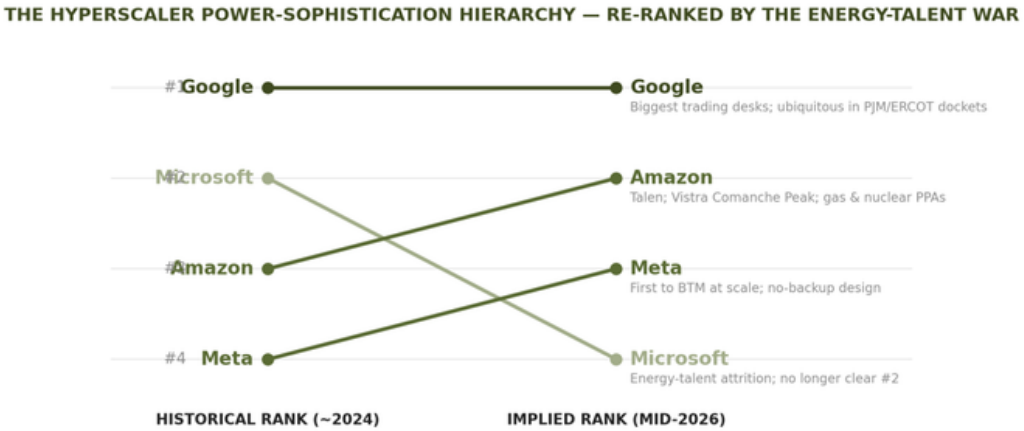

Source: Keel synthesis of Catalyst discussion (Kann × Eliahou Ontiveros, July 2026). Rankings are the participants' qualitative judgments, not a published index.

Google remains the clear No. 1 — and, per Ontiveros, not merely ahead but structurally different: the largest energy trading operation among the four, flexibility deals at the frontier of load-management practice, a name that appears in ISO stakeholder minutes more than any peer's, and an energy team credibly a multiple of the size of the next hyperscaler's.

Microsoft's slide is the news. Kann, who long ranked the company a clear second, now calls the ranking genuinely unclear after a series of senior energy departures; Ontiveros is blunter and simply ranks it fourth. **Amazon has stepped up dramatically** — Talen, Vistra Comanche Peak, and large-scale PPAs spanning gas, nuclear and renewables. **Meta is doing "the most unusual stuff":** first mover on behind-the-meter at scale and, as a first-party company with no cloud tenants demanding four-nines, willing to run modern data centers with essentially no backup — on the order of 10 MW of gensets against 100 MW of load. The Columbus complex is the emblem: tent-structure data centers fed by a medium-voltage line between two existing substations, curtailment-exposed on peak days, as a bridge while on-site plants are built. Google and Amazon have largely avoided BTM by being the most aggressive depositors across the U.S. utility landscape; the ~150 GW of contracted-load queues utilities reported at end-2025 — are, in his read, mostly those two names.

THE LABS SET THE LOAD CURVE

The most important reframing of the episode: to understand power demand today, the focus should not be on the hyperscalers but rather the AI labs. OpenAI and Anthropic are now the 'demand' — roughly half the combined gigawatts Amazon and Microsoft stand up each quarter are built for the two labs, making the clouds increasingly proxies for lab buildouts. The labs also work the problem directly, evaluating powered-land sites themselves before bringing in partners once a site clears. Their binding constraint is credit, not intent: neither is investment grade, and at gigawatt scale the capex runs $50 per GW billion self-built or $60–70 billion contracted.

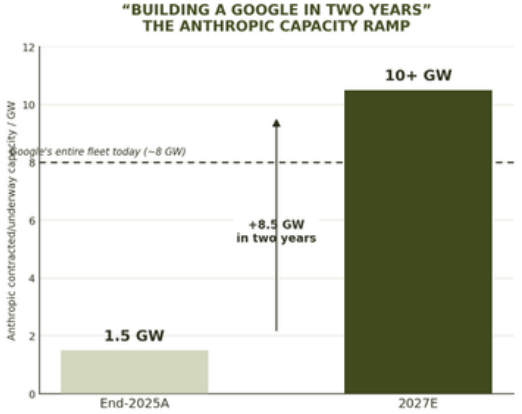

> *"If you're an AI lab, power is the lifeblood of your business. Power is revenue… You cannot make that kind of investment decision if there's uncertainty on the timeline… Anthropic: 1.5GWs of capacity end of 2025; by '27 they want over 10 GWs… That's the size of Google today… That's not going to happen with the grid. There's just no way… If you want to be in control of your destiny, the best option from a buyer's perspective is to simply bring your own generation."*
>
> **— Jeremy Eliahou Ontiveros, SemiAnalysis, Catalyst interview (July 9, 2026)**

Catalyst discussion: Reflects the targets for power additions that Anthropic has announced

The divergence worth logging sits on the supply side. Kann bets the U.S. will add substantially more generation over the next 12 months than SemiAnalysis projects — in his analysis the capacity constraint is real but overstated, while the transmission-and-distribution constraint is as bad as feared or worse. Ontiveros holds that both bind. The distinction matters commercially: if Kann is right, the scarce late-decade asset is not generation per se but deliverability — existing interconnection positions, substations and wires. Either way, both land on the same verdict: even optimistic firm-equivalent additions look like rookie numbers against a compounding load book, and self-supplied generation converts an unbounded timing risk into an execution risk the buyer controls.

SILICON WARS: BACKSTOPS AS POWER FINANCING TOOL

Kann and Ontiveros frame the whole power race against a background of 'Silicon Wars.' Google struck first — and it did so with a financing instrument rather than a chip. Understanding that its largest customer, Anthropic, would inevitably go vertical, and seeing the opening to sell TPUs against Nvidia, Google began guaranteeing third-party data center buildouts serving Anthropic. The TeraWulf, Hut 8 and Cipher transactions with Fluidstack — roughly a gigawatt in aggregate — set the template: Google's credit signature stands behind the sites if the off-taker fails, which is precisely what makes the construction financeable. The program has scaled to roughly $50 billion of balance-sheet obligations, call it 4–5 GW of capacity, on which Anthropic buys TPUs at approximately $20 billion per gigawatt: $50 billion of contingent support generating on the order of $100 billion of hardware revenue. It is the Nvidia-neocloud playbook executed an order of magnitude more aggressively — Nvidia bought neocloud equity and left the financing burden on its partners' books; Google is directly underwriting the buildout.

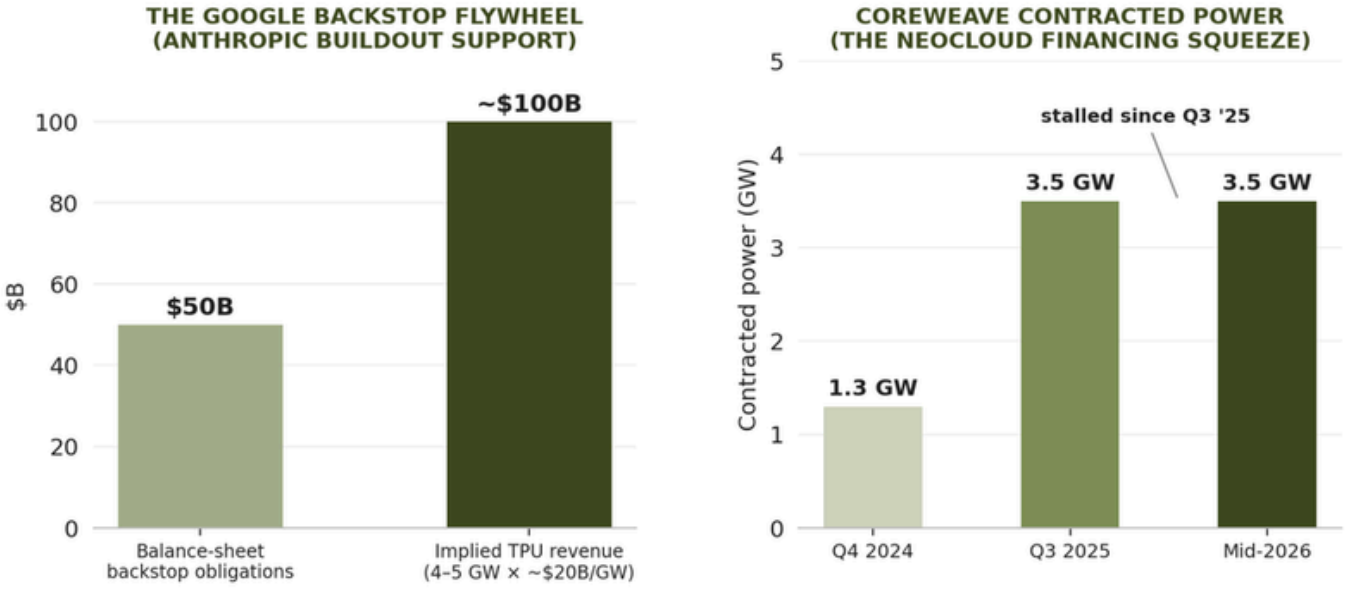

Sources: Catalyst discussion. Left: ~$50B obligations ≈ 4–5 GW; TPU pricing ~$20B/GW per Eliahou Ontiveros. Right: CoreWeave contracted power as cited (1.3 GW Q4 2024 → ~3.5 GW, flat since Q3 2025).

The mirror image is the neocloud financing squeeze. CoreWeave's contracted power ran from ~1.3 GW at the end of 2024 to ~3.5 GW today — and has slowed since Q3 2025, coincident with the high-yield selloff that froze the market for exactly this paper. Meanwhile the physical market has turned more challenging: turbine deposits, swelling large-load interconnection deposits, utilities demanding multi-billion-dollar commitments per gigawatt. Non-investment-grade builders are structurally disadvantaged at precisely the moment scale demands credit. Unsupported GPU capacity now trades like technology risk; hyperscaler-backed energized capacity trades like infrastructure credit — and the winner is not just the model lab or the chip vendor, but the platform that can make power bankable.

The forward call is Nvidia's response. Every gigawatt Google underwrites is TPU share taken from GPUs — an arrangement Ontiveros describes as approaching existential for Nvidia — and he expects major Nvidia balance-sheet-support announcements to unlock a second wave of neocloud growth in 2H2026. If that materializes, the universe of creditworthy off-takers for developer sites expands materially.

SPEED TO POWER SORTS THE EQUIPMENT STACK

Kann and Ontiveros wrap up their discussion by highlighting how the equipment market now organizes on a single axis — time to energization — in a cascade running from H-class turbines at the premium end down through aeroderivatives, reciprocating engines and fuel cells. xAI set the playbook in 2024 with aeroderivative units in 20–40 MW blocks; OpenAI (via Crusoe and Oracle at Abilene) and Meta (Columbus) followed in 2025. The frontier of scale is now reciprocating engines: the OpenAI/Oracle project in Shackelford County, Texas will deploy 2.3 GW of ~4 MW recips with flywheels handling GPU load transients, and both high-speed and medium-speed units are being adopted at gigawatt scale through the leasing channel — even Meta and Microsoft are shifting from self-build to leasing from operators who supply BTM power.

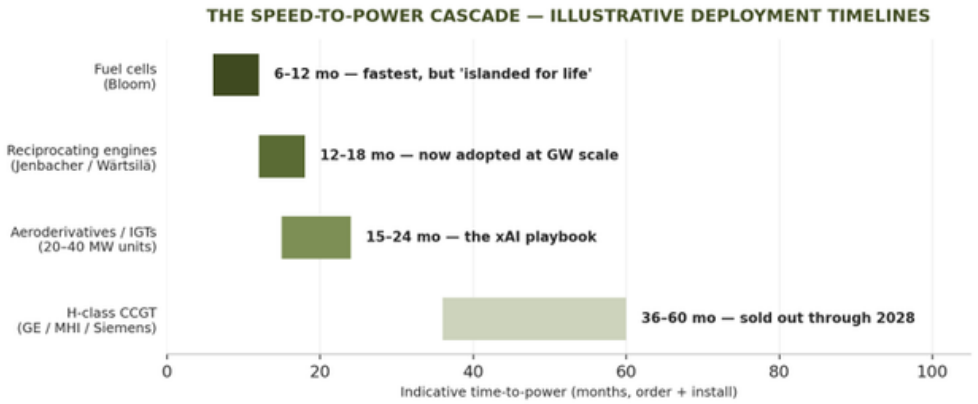

Illustrative timelines synthesized from the Catalyst discussion. Indicative, not vendor quotes. SemiAnalysis calls a 2026 peak in gas-turbine orders as utility buying fades.

Two structural notes complete their map. West Texas is the inflection: both participants expect the end-state gigacampus — tens of thousands of acres combining gas, solar, batteries and possibly wind — to concentrate there, with Lancium's land-and-power position behind Abilene and Crusoe's wind-coupled Armstrong County site as early markers, and credible developers now planning 5 GW campuses where vaporware once stood. ***The Information's*** reported last week that Blackstone-backed Lancium's portfolio of gigawatt-scale Texas campuses has drawn interest from several tech companies, including Nvidia. And clean firm is optionality, not supply: the nuclear and geothermal announcements are real but largely non-binding and milestone-contingent — options on late-decade capacity at best, deprioritized today against anything that delivers speed to power.

Bottom Line: the equipment cascade is why BTM is not one technology. It is a timing strategy. Developers are arbitraging whatever equipment can deliver firm electrons inside the AI deployment window.

MSL Y1: Meta INNOVATES WITH BTM GAS & DISTRIBUTED RL

If the Catalyst interview explains why power is revenue, Meta — one year into the Superintelligence Labs rebuild — shows what the strategy looks like in steel, gas turbines, fiber and project finance. The product evidence remains unclear; the input evidence is overwhelming, and for Digital Power it is the inputs that matter. SemiAnalysis did a deep dive on it all and judges Meta the only hyperscaler on track to be world-class in all three frontier inputs — data, talent and compute — and projects it will end 2026 holding more AI compute than either OpenAI or Anthropic.

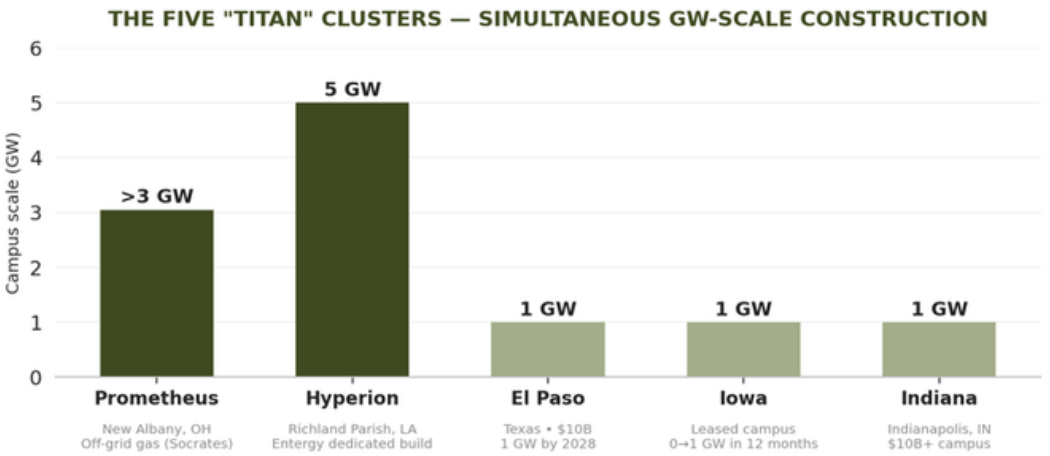

The compute ramp is real and externally verified. No full 1 GW campus had ever been under construction before; Meta currently has two. Prometheus has expanded from ~1 GW to more than 3 GW in two years as a constellation of 27 data centers across six campuses, partially operational. Hyperion is sized for ~5 GW of IT load on 2,700 acres, with 1.5 GW under construction including three 400 MW single buildings — the largest ever built. El Paso was upsized in March from $1.5 billion to $10 billion and confirmed at 1 GW by 2028; Iowa went from greenfield to full-gigawatt construction in twelve months; Indiana is a $10 billion-plus campus near Indianapolis.

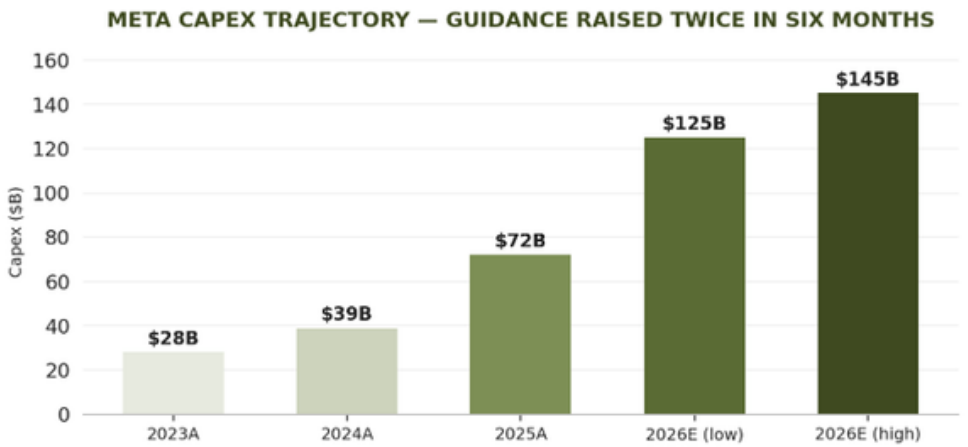

Sources: SemiAnalysis Datacenter Model; Meta earnings releases and guidance; Entergy; OPSB filings. 2026 range guided at Q1 earnings.

The capital follows the concrete. Meta guided 2026 capex to $115–135 billion at Q4 earnings, then raised to $125–$145 billion at Q1 — nearly double 2025 and more than 2024 and 2025 combined. Multi-year infrastructure commitments grew $107 billion in the first quarter alone. Zuckerberg's stated willingness to run free-cash-flow negative, plus the absence of a cloud-resale business competing for internal compute, underwrites the projection that Meta ends the year with the largest AI fleet of any lab.

THE LARGEST VALIDATION of BEHIND-THE-METER GAS-TO-POWER

SemiAnalysis notes that "some of these facilities use behind-the-meter power." The full record is far stronger, and it is one of the most strategically significant elements of the report. Meta is running three distinct power procurement models simultaneously — a deliberate architecture built to escape queue risk and preserve deployment control.

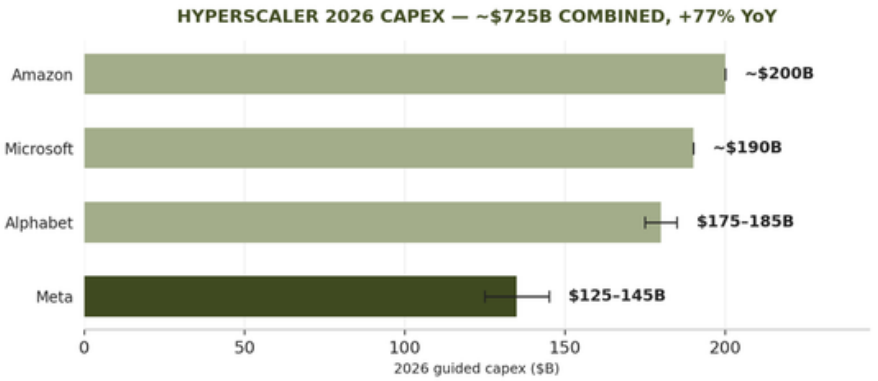

Sources: company guidance per earnings calls; aggregate ~$725B, +77% vs. ~$410B in 2025.

MODEL	SITE	STRUCTURE	KEY FACTS
Pure BTM / off-grid	Prometheus, OH	Williams (Will-Power) builds/operates; Meta affiliate offtakes under 10-yr agreement	2 × 200 MW plants (Socrates N & S), ~$1.6B; 30-unit modular fleet; two 24-in gas laterals; no PJM interconnection by OPSB condition
Dedicated utility build	Hyperion, LA	Entergy Louisiana constructs; Meta pays full cost of service; LPSC-approved	10 gas plants / 7.5 GW / ~$11B; +2.5 GW renewables & storage; 240 mi transmission; ~30% addition to LA grid capacity
On-site bridge gen	El Paso, TX	813 modular gas generators owned at site; bridging grid delivery	366 MW, ~$473M, online 2027 — ahead of grid service for the 1 GW campus

The Ohio structure deserves particular attention. Ohio's HB 15 (2025) carved self-generators and mercantile self-power systems out of the public-utility definition, explicitly permitting third-party ownership and PPAs — the enabling framework for Socrates. The plants are off-grid by regulatory condition: they do not participate in PJM wholesale markets and avoid the multi-year interconnection queue entirely. Industry BTM data-center capacity stands at roughly 2 GW today and is projected to reach ~13 GW by end-2027.

The countervailing signal arrived in March: Senate Democrats opened a probe into gas-powered AI data centers, questioning Meta directly on New Albany. Federal attention on the BTM gas model is now live. BTM solves the queue; it does not eliminate social-license risk — and permitting posture and ratepayer optics will govern how quickly the model scales from here.

DISTRIBUTED RL BREAKS THE MEGA-SITE CONSTRAINT

One of the report's most interesting conclusions is the emergence of Distributed RL as a development vector. The technical architecture matters for power siting. Physics forces the design: latency within a data center cluster runs 1–10 microseconds, but a site 100 kilometers away cannot be reached in under ~500 — light in fiber sets the floor. Synchronous pretraining therefore stays within one region; RL post-training, whose rollout generation is very parallel and staleness-tolerant, could potentially be spread out in smaller data centers. Currently, Meta's AI-Backbone will connect Titan campuses up to 2,000 kilometers apart, and the effort is not just limited to them: Google DeepMind's Decoupled DiLoCo cut required inter-datacenter bandwidth from ~198 Gbps to 0.84 Gbps across eight sites, and Microsoft's Fairwater AI WAN explicitly disaggregates the training pipeline by latency profile. When Meta, Google and Microsoft converge on the same architecture within a single quarter, it's likely more than a hypothesis.

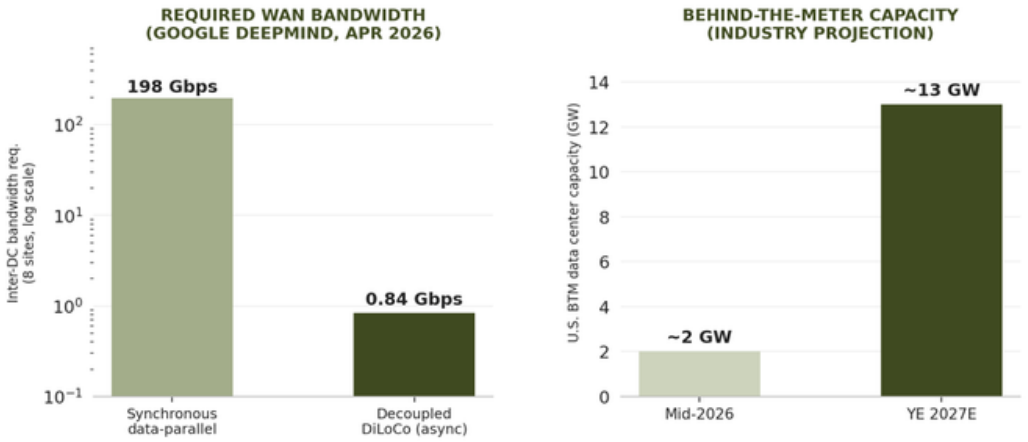

Left: Google DeepMind, "Decoupled DiLoCo" (Apr 23, 2026) — inter-datacenter bandwidth for an 8-site training run. Right: industry BTM capacity projection per trade reporting.

The data center siting consequence is not insignificant: a growing share of frontier training FLOPs could no longer require the single contiguous gigawatt. In this situation, selection of this criteria for RL and rollout capacity collapse to speed-to-power, cost of power, and fiber access — axes on which a 50–500 MW site is fully compatible with megasites in a single portfolio. The tenant universe widens accordingly: any lab running distributed RL is a potential offtaker for mid-scale capacity, expanding the universe beyond the hyperscalers who have been focused on megasites. Megawatt quality still matters; the form factor and geography of where all of it must sit may become far more flexible.

It's worth nothing that RL's raw materials — tasks, environments, verifiers, recorded work trajectories — have spawned a supply chain from nothing: the three incumbents are all at $1 billion-plus ARR within two years, frontier labs pay $5,000 or more per quality coding task, and top contractors earn seven figures. Meta's response has been characteristically blunt: workplace tracking software to capture the scarcest RL input, real recordings of white-collar work, and the reassignment of thousands of engineers to full-time task and environment creation. The asset is real and probably decisive if retained; contemporaneous reporting of involuntary transfers, internal petitions and elevated attrition says the execution risk is not immaterial.

AI CAPITAL BECOMES INSTITUTIONAL CREDIT

The third major theme of the week is the maturation of the financing stack. AI infrastructure is no longer funded only by hyperscaler cash and venture equity; it is increasingly packaged into bonds, 144A private placements, leveraged loans, joint ventures and stabilized real-estate vehicles. The 144A rail — a $3.5 trillion institutional market hidden in plain sight — has absorbed $71.9 billion of AI infrastructure paper since May 2025, a mechanism especially valuable to former bitcoin miners and neocloud developers without long bank-construction track records. What lenders are underwriting reduces to four questions: tenant credit/tenor, power firmness, GPU residuals, and local legitimacy. Projects that solve all four finance first — and the cost-of-capital gap compounds, because they can then outbid everyone else for power, equipment and land.

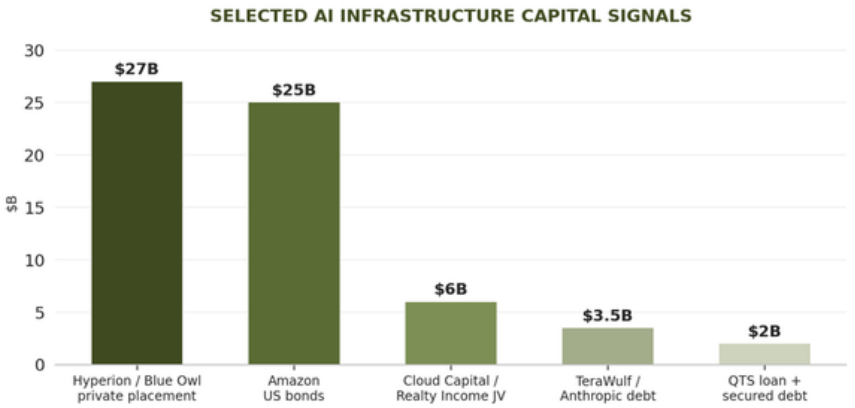

Sources: Bloomberg, The Information, Keel synthesis. Selected transactions are not directly comparable but show the capital rails forming around AI infrastructure.

GAS IS THE BRIDGE - AND THE UNDERWRITING RISK

The last major theme this week came from a Wood Mackenzie report: cheap Henry Hub gas made U.S. data-center power globally advantaged, but AI load, LNG exports and maturing supply are pushing the benchmark toward roughly $5/MMBtu real by 2035, with power-sector gas demand forecast to rise by up to 17 Bcf/d and U.S. LNG capacity more than doubling. Gas remains the only dispatchable bridge that can scale close to AI deployment timelines — and it grows more valuable as renewables push balancing demand higher. But long-dated data-center PPAs now carry commodity exposure that cannot be assumed away: fuel strategy has become part of site selection, financing and tenant pricing, and the best projects will combine gas access, hedging, grid optionality and a credible clean-power addition path.

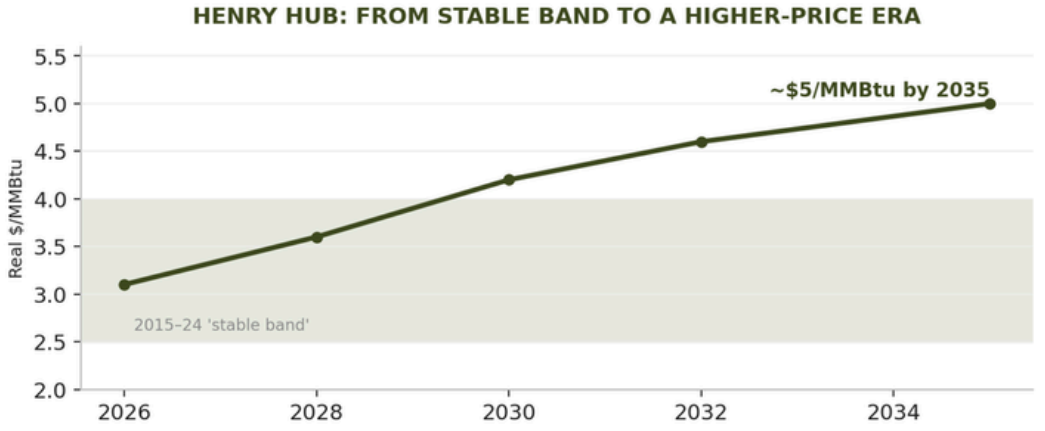

WEEKLY ROUNDUP

FIFTEEN STORIES, ONE FLYWHEEL

The week's top stories, itemized

1. Catalyst × SemiAnalysis: power strategy re-ranks hyperscalers; energy talent becomes a battleground.
2. Meta MSL Year One: five Titan clusters, BTM gas, and distributed RL convert compute ambition into megawatts.
3. 144A market: AI borrowers tap a $3.5T institutional rail hidden in plain sight. The rail offers liquidity-like access to institutional investors without full public registration — especially valuable to former miners and neocloud developers that lack long bank-construction track records — and has absorbed $71.9B of AI infrastructure paper since May 2025.
4. TeraWulf / Anthropic: $3.5B debt planned against a 20-year, ~$19B Kentucky lease. Project vehicles are now borrowing directly against the cash flows of the ultimate AI-lab tenant; a miner-conversion campus underwritten on lease economics is the clearest proof yet that these assets can be financed as infrastructure rather than as technology risk.
5. Big Tech debt / Amazon bonds: hyperscaler balance sheets become infrastructure funding vehicles. Amazon's ~$25B bond print showed institutional appetite for long-duration AI power cash flows at mega-cap scale; the hyperscaler balance sheet is becoming the cheapest funding rail in the stack, and its spread advantage compounds into land, turbines and talent.
6. QTS, Csquare, Realty Income, Segro: data centers complete the crossing into core real estate and structured credit. Four transactions in nine days trace the full institutional capital stack. At the top, Realty Income — The Monthly Dividend Company, the definitional core-net-lease REIT — committed up to $1.4B for a 45% stake in a Northern Virginia hyperscale portfolio through a $6B+ programmatic JV with Cloud Capital and a global institutional investor: stabilized assets, investment-grade tenants, 15–20 year triple-net leases, with a European expansion mandate. In the credit stack, Blackstone's QTS closed an $870M secured package against five campuses in Phoenix, Richmond, and DFW — the latest draw on a securitization program that already includes a $3.5B CMBS, in a market where data center ABS/CMBS issuance doubled year-over-year and the asset class now represents 13% of the entire single-asset single-borrower market. In public equity, Brookfield filed Csquare's $1.35B IPO at a $4.2B valuation — with proceeds earmarked not for growth but to retire a $734M revolver and a slice of $4.3B in asset-backed loans, public shareholders refinancing the securitized layer. And in Europe, Segro's second JV with Pure DC — a 48MW, ~£0.8B Paris build in which Segro contributes powered land carrying 75 MVA from its 3.0 GVA European power bank against just £60M of cash equity — arrives mid-takeover-battle, with Prologis's £12.6B bid contested in part over who is better positioned to monetize that data center pipeline.
7. Wood Mackenzie Henry Hub: AI and LNG push gas toward a higher-price era. The forecaster sees Henry Hub migrating toward ~$5/MMBtu real by 2035 as power-sector demand adds ~17 Bcf/d and U.S. LNG capacity more than doubles. The molecule that solves the 2027 power problem becomes the 2035 margin problem; fuel basis now belongs in the underwriting model.
8. Anthropic IPO financials: profitable lab economics keep bidding for compute and power. The disclosures quantify what the ramp implied: a lab running from 1.5 GW of contracted capacity at end-2025 toward 10+ GW by 2027, with 90%+ of revenue B2B and more than 70% of ARR in coding — demand durable enough that Google has placed ~$50B of balance-sheet backstops behind the buildout and sells TPUs into it at roughly $20B per gigawatt. The load book is bankable, which is precisely what lets power-backed paper price like infrastructure.
9. Rise of the Neoscalers: A five-rail capital stack has hardened into place — core-infrastructure and ABS takeout at the top, senior project finance at 60–70% loan-to-cost, then 144A high-yield, leveraged loans, and GPU-backed mezzanine at the bottom — and the sorting is already visible in a single name. CoreWeave built its platform on the bottom rails: contracted power ran from 1.3 GW at end-2024 to 3.1 GW at end-2025, funded with high-yield 144A paper priced for a single-counterparty credit story. Then the counterparty mix crossed the quality threshold. In Q1 2026, with contracted power topping 3.5 GW, backlog at $99.4 billion, and non-investment-grade customer concentration falling below 30%, the company executed an $8.5 billion investment-grade, HPC-backed debt facility — its first — and lowered its weighted average cost of debt in the same quarter it signed Meta and Anthropic. The steel didn't change. The contracts did. The physical platform was always durable; the contract decides who refinances, at what rating, and on which rail.
10. Meta Canada: Meta's first Canadian data center — a C$13 billion, 1 GW campus in Sturgeon County, scalable to 1.8 GW and the company's largest outside the United States — extends the supercluster program (Prometheus in Ohio coming online this year, Hyperion in Louisiana scaling toward 5 GW, and the additional "titan clusters" Zuckerberg has promised at El Paso and Beaver Dam) beyond the $600 billion Meta has committed to U.S. infrastructure through 2028. And it was chosen the same way: gas-backed firm supply first, everything else second. The anchor is a power purchase agreement with the Greenlight Electricity Centre, a 932 MW gas plant backed by Pembina Pipeline, Kineticor, and Morgan Stanley Infrastructure Partners, which reached FID six days before Meta went public — with Capital Power's existing gas fleet bridging supply until Greenlight enters service in late 2030. At full build the campus consumes roughly 150 MMcf/d, and Meta says it worked with the AESO, AltaLink, and Capital Power to plan its energy needs years in advance of the announcement. The site selection is the Henry Hub thesis in miniature: cheap Alberta gas at a discount to Henry Hub, a regulatory framework built to let large loads bring their own generation, and dry cooling in a cold climate.

WEEKLY ROUNDUP

11. Microsoft emissions: AI buildout pushes Microsoft's Scope 3 purchased electricity (market-based) emissions to roughly 20 million metric tons of CO_2e. The figure lands against approximately $190 billion of 2026 capex and a power organization that the Catalyst interview ranked fourth among the hyperscalers following a series of senior energy leadership departures. The increase is being driven primarily by the emissions associated with electricity purchased to power Microsoft's expanding AI infrastructure—the buildout's carbon bill arrives through the grid rather than from Microsoft's own direct fuel combustion. That sharpens the tension between the company's climate commitments and its growing need for firm power, while giving critics of the natural gas bridge a concrete emissions benchmark to target.

12. GOP candidates break with Administration: data-center backlash becomes bipartisan electoral terrain. The defections follow the tape: Blackstone's QTS walked away from Virginia's ~$100B Digital Gateway this month, Carbon Direct counts roughly $170B of U.S. capacity blocked, stalled or canceled since 2024, and data centers are flagged as a genuine swing issue in Virginia and Utah ahead of the midterms. Ratepayer and land-use grievances no longer map onto party lines — a material repricing of permitting risk for every large load.

13. NYT influence reporting: The battle over social license is no longer confined to county planning boards and zoning hearings. Recent reporting in The New York Times shows that Chinese and Russian influence operations have attempted to amplify existing community opposition to AI infrastructure through coordinated online campaigns, using AI-generated images, fabricated local news content, satellite imagery, and social-media messaging focused on electricity prices, water consumption, noise, and environmental impacts. OpenAI identified and disrupted one China-linked operation —dubbed "Data Center Bandwagon"—that used ChatGPT to generate anti-data-center content designed to exploit legitimate local concerns and deepen public distrust. The New York Times also documented Chinese state media and Russian influence networks seeking to inflame debate around AI infrastructure as part of a broader geopolitical competition over U.S. technological leadership. Importantly, neither OpenAI nor independent researchers found evidence that these campaigns materially shifted public opinion or drove permitting outcomes; rather, they sought to amplify authentic local grievances that already existed.

14. Meta Cheyenne water: The Cheyenne Board of Public Utilities stopped accepting industrial wastewater from data center fill-and-flush and closed-loop cooling operations after tracing a rare bacterium in the city's reclaimed water to Goat Systems LLC, the entity Meta uses to build its Cheyenne campus. The Board found subcontractor Goat Systems in significant noncompliance for discharging water carrying Cupriavidus gilardii, a metal-resistant bacterium that interfered with two water reclamation plants and pushed the reuse system offline for months of cleanup. Fill-and-flush privileges were revoked March 24, and a wider suspension now covers every data center connected to city services. The irony is the marketing angle: hyperscalers market sealed liquid loops as a near-zero-water alternative to evaporative cooling — but the one-time fill is the step that produces a discharge, and the flush water leaves the site before the loop is sealed. Cheyenne's concern extends beyond the bacterium, because closed-loop systems can carry glycol and other chemicals municipal treatment plants aren't built to process, and the city sprays reclaimed water on parks and golf courses.

15. HBM / DRAM scarcity: SK Hynix's $26.5B US debut Friday — the largest foreign listing in American history, surpassing Alibaba's 2014 record, and the second-largest US listing ever behind SpaceX four weeks prior — priced the memory shortage as an asset class. The ADRs, seven times oversubscribed, closed their first session up 13%; the underlying Korean shares have returned 634% in twelve months, carrying the company past a $1 trillion valuation and past Samsung as Korea's most valuable listed firm. The market is paying for position in a three-supplier complex in which SK Hynix alone controls 56.4% of HBM, and in which scarcity has already crossed into hyperscaler guidance: Meta named memory prices explicitly in raising 2026 capex to $125–145B. Every powered megawatt now carries a memory surcharge the power layer cannot hedge.

The supply response is equally instructive. Seoul's answer to the shortage — ₩1,350 trillion (~$880B) in coordinated public-private commitments, anchored by the $518B "K-Semiconductor Grand Leap" that puts four new Samsung and SK Hynix mega-fabs in the Honam region, plus $356B of AI data centers through 2035 — is the upstream mirror of the American power buildout, and the siting rationale is the tell. The fabs are going to the southwest for one stated reason: the Seoul cluster has exhausted its utility capacity, and Honam still has surplus power and water. Samsung cited power and water incentives by name in selecting Gwangju; Korean brokerages are already arguing the real trade is the grid and water infrastructure, not the silicon. Load-follows-power has graduated from developer heuristic to state industrial policy. The caveat cuts in favor of the thesis: even on schedules accelerated by twelve years, the new fabs don't produce until the mid-2030s, which is why the shortage — and the pass-through into compute economics — is structural through the decade. IPO proceeds fund Yongin fab capacity and Cheongju advanced packaging; Commerce Secretary Lutnick is simultaneously pressing both Korean makers to build in the US, with Micron countering at $250B of domestic commitments.