THE AI STACK WEEKLY ISSUE 29 · WEEK 45 OF 2026

THE BOTTOM LINE

Who pays the rent if the lab cannot: six financings sorted AI compute by guarantor, from Meta's 6.6% to a vendor convertible that costs 11-15% as equity

The house measurement this week prices Anthropic's $42B Broadcom rent facility as capital: 11.3% a year held to its five-year maturity, 14.9% a year if the holder converts after two years. The ladder it sits on sorts the week's financings by who stands behind the rent, not by what the chips earn; the sort is borrowed and credited: Haircut ('Who actually pays if the AI data centre bet goes wrong', nodiligence.substack.com, October 7), UBP (September 16) and a September 2026 single-campus bond survey read through Clifford Chance's three-bucket framework, via Luminix, each tiered data-centre debt by obligor (the party bound to pay); all are outside the graded packet. What this issue adds is the price of the facility and the placement of six in-window events on one ladder: two bond prints, one equity print, one guarantee signed, one facility's terms disclosed and one lease entering syndication, plus two market prices (CoreWeave's bonds and its credit default swaps). Meta sold $15B of senior notes on November 3 at a size-weighted coupon of 6.60% (Treasuries +110 to +220 bp, a basis point being a hundredth of a percentage point; rated Aa3 stable / AA- negative) on its own name, with no project, lease or vendor in the structure. Oracle, two letter grades lower at Baa2/BBB-, filled its Investor Day cap on November 2 with $10B of notes at 7.77% size-weighted (Treasuries +240 to +335 bp) and $10B of stock at $116, an 11.2% discount to the last close. Anthropic, unrated and pre-IPO, has rent on two leases that someone else stands behind, in two forms: a guarantee and a vendor loan. Google signed credit-support agreements (TeraWulf 8-K, November 4) under which, if Anthropic defaults on the ~401 MW, ~$19B lease, Google assumes it or pays an amortising termination fee, until Anthropic has held investment-grade ratings for three months. The S-1/A

(amended IPO registration statement) disclosed that Anthropic can fund quarterly rent on its $125.2B Broadcom TPU lease by issuing up to $42B of notes to Broadcom at 9% paid in kind (PIK: interest added to principal rather than paid in cash), convertible at the holder's election at 90% of the trailing volume-weighted average price (VWAP). SpaceX, a strong obligor with no guarantor, is leasing ~$47.3B of Rubin equipment (about $40B of financed principal plus a ~$7B lease prepayment; Bloomberg, grade 3) as an absolute obligation to pay rent with no residual-value guarantee, and its $8.3B senior-secured loan began pre-marketing on November 5. CoreWeave, a levered operator with no guarantor and no investment-grade parent, saw its 9.625% 2032 notes yield above 16% and its five-year CDS (the annual cost of insuring its debt) pass 1,289 bp on November 5 (Deutsche Bank via CNBC).

The arithmetic: a note held one year accretes to $1.09 and converts into $1.21 of stock, 21.1% all-in; held two years, 14.9% a year; held to maturity, 11.3%; under the penalty terms that apply when revenue conditions are missed (13% PIK, 85% conversion), 16.7% at maturity. These are costs of issuing equity through a note, not borrowing costs, and the short-dated figures depend on Broadcom choosing to convert early (Anthropic has no call), so the like-for-like comparison on the week's tape is Oracle's 11.2% discount on $10B of stock, which the facility's cheapest path matches and its penalty path exceeds by half. The facility is also small against the rent it exists to pay: $42B covers 33.5% of the $125.2B of principal rent, 6.7 quarters at an even run-rate, a bridge on one lease rather than a financing of it. The frames on the week's other events are credited in the changelog.

What changes this week. For a board or CFO at any company carrying vendor-attached 'financing capacity', book it as forward equity sold at a standing discount, not as liquidity, and put it last in the waterfall: a facility that costs 11-15% as equity belongs behind Oracle-grade unsecured debt. For an IPO investor in the lab, model the $125.2B lease and the $69.8B of commitments due in twelve months, not the $42B facility; the dilution overhang (5-7% if fully drawn and held) triggers when the company is most stretched. For an operator or architect buying capacity from a neocloud, ask who stands behind the landlord's rent and whether that party has silicon conditions attached; a backstop conditioned on one vendor's chips is a procurement constraint on you as much as a credit enhancement for them. For anyone underwriting the 2028 pipeline, Nomura's point stands and the ballot extends it: 40 of 41 local measures favoured limits, Texas's governor-elect has pledged an executive pause pending legislation, and the Georgia PSC (Public

Service Commission, the utility regulator) flipped; the rung that just repriced is the one that needs new permits.

The ladder has two rungs the market has not yet priced: the lab whose rent is guaranteed by an AA+ hyperscaler, because the TeraWulf project debt is not yet raised (TeraWulf sought ~$3.5B in July per DCD), and the lab funding rent through a vendor's convertible, because the facility is undrawn. The November 5 equity moves should not be read as the ladder: hyperscalers fell 2-4%, neoclouds 10-17%, and semiconductors and memory, which are not on the ladder, 9-22%, a sort that looks like beta rather than guarantor, though one day's tape cannot settle that. IFR's line that data-centre financing has 'ground to a halt', with five IPOs pulled against roughly $330B raised this year (about $100B of it debt), is true of the levered unguaranteed rung and noise for the market: $25B of hyperscaler bonds and $10B of equity cleared in the same three days.

Beyond the week, the direction of the guarantees matters: sellers of compute are lending their credit to buyers of compute so that a third party will finance the buyer. NVIDIA's August 17 8-K (outside the graded packet) disclosed the vendor-side instance first: residual-value guaranties on OpenAI's SB Energy leases, capped at $105B, with NVIDIA as exclusive compute provider. Google's TeraWulf agreement is the hyperscaler-side instance and carries the same tell: Exhibit 10.1 to the 8-K, read after the research stage, requires Google-approved TPUs at the premises for the term (s.2(k)) and limits Google's option to pay the termination fee instead of assuming the lease for six years (s.4(e)). Mosaic Singularity and alfw.ai read the backstop as a silicon distribution contract written as a credit instrument on November 5, before this issue; this issue adopts the reading and adds the vendor-side pairing. The framework's second hypothesis (capital concentrated, returns diffuse) now extends to the credit itself: a lab's ability to rent compute at a financeable rate depends on hyperscaler and vendor balance sheets.

The fundamentals are not in dispute: AMD Data Center $8.55B (+97%), Arista +58.5% with a ~$18.5B 2027 outlook, Coherent with no cancellations, Eaton orders +101%, and the buyer printed $16.96B of revenue against a $567M operating loss that includes a $2.0B non-cash charge for its donation-matching programme (about +$1.4B and an 8.5% margin without it, level with Q2's 8.4% even as compute rose from 48% to 58% of revenue). The revenue is real and growing, the rent is larger and growing faster, and the gap is being bridged by guarantees from the parties that sell the compute.

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

November 7, 2026

Issue 29

THIS WEEK IN 30 SECONDS

Executive summary.

- Anthropic's $42B Broadcom rent facility, priced as capital, costs 11.3% a year held to maturity and 14.9% if the holder converts after two years (house measurement); those are costs of issuing equity through a note, so the like-for-like comparison is Oracle's 11.2% discount on $10B of stock this week, not any coupon. Around it the week's financings sorted by who stands behind the rent: Meta 6.60% on its own name, Oracle 7.77% plus the equity discount, Google guaranteeing Anthropic's ~$19B TeraWulf lease until the lab is investment grade, and CoreWeave, with no guarantor, above 16%. Book vendor-attached financing capacity as forward equity, behind unsecured debt in the waterfall.
- The November 5 selloff (Nasdaq -5.5%, SOX -12.3%, where SOX is the Philadelphia Semiconductor Index; Nvidia -9.2%) came two days after Democrats took both chambers and data-centre-skeptic governors won Texas, Ohio and Georgia, while the same week's primary-source prints were uniformly strong (AMD Data Center +97%, Arista +58.5%, Coherent 'no cancellations'); the hardware lens's reading is that the market cut 2028-plus expectations, not 2026-27 shipments, so hold acceptance tests to delivered racks rather than renegotiating allocation.
- Marvell signed the definitive all-stock agreement for Lumentum at a $115B implied enterprise value (4.500 shares, 23.9% premium, close H1 2028), resolving last week's p148 as a hit seven weeks early. The premium is ordinary for public M&A; the pricing-power evidence is the order books (Coherent: no cancellations, coverage into 2028), and the networking hypothesis stays strained until the operator-layer test (p154) resolves. Get allocation and change-of-control language into optical contracts before the close.
- Cyber capability became the distribution gate, and the labs gate on different quantities: OpenAI shipped GPT-6 Cyber, rated Critical (its top internal risk tier), only to vetted Daybreak Red organisations (its trusted-defender programme) a day after disclosing, by its own count, 948 distinct Linux local-privilege-escalation (LPE) vulnerabilities with working exploits for $1.21M; Cognition withheld SWE-2 Flash weights over an uplift on its own ExploitBench (Cognition-run, 11.3% to 34.8%) while serving models that score higher; Gemini 4 Carbon stayed API-less at 60. Moonshot published Kimi K3.1 Pro's weights on schedule at 46, the first past 40 with public weights on v5. Budget defender tooling against the most capable served model.
- Anthropic's S-1/A (the amended IPO registration statement) put a $16.96B quarter (+47% QoQ) next to a $567M operating loss, of which $2.0B is a non-cash donation-matching charge; excluding it, operating income was about +$1.4B, an 8.5% margin level with Q2's 8.4% even as compute rose from 48% to 58% of revenue. Commitments total $895.4B ($69.8B due in 12 months) against $146.6B of liquidity. METR measured the lab improving its models ~1.9x faster than before and estimated ~80% of that as AI-assisted R&D, three days after the filing said it 'may imminently' meet its own acceleration threshold.
- Power printed as a supplier's own constraint for the first time in this series: Applied Optoelectronics cut its end-2026 transceiver capacity target to >500k/month from ~650k on electrical-power availability and construction delays; Eaton tracks 366 GW of announced data-

centre projects against ~50 GW installed; Arm cited Morgan Stanley's 38 GW global power-shortfall projection through 2028; and Google answered with $17B in Wyoming on 702 MW of fuel cells. Assume permitting, not silicon, gates 2028 capacity in Texas, Ohio and Georgia.

BY THE NUMBERS

-12.3%

Philadelphia Semiconductor Index on November 5, worst day since March 2020

$115B

Marvell-Lumentum implied enterprise value, definitive all-stock agreement

11.3-14.9%

Annual cost of Anthropic's $42B Broadcom rent facility as equity issuance, held to maturity or converted at two years (house measurement; Oracle sold stock at an 11.2% discount the same week)

+97%

AMD Data Center revenue year on year, $8.55B, Q3 2026

>16%

Yield on CoreWeave's 9.625% 2032 unsecured notes, November 6

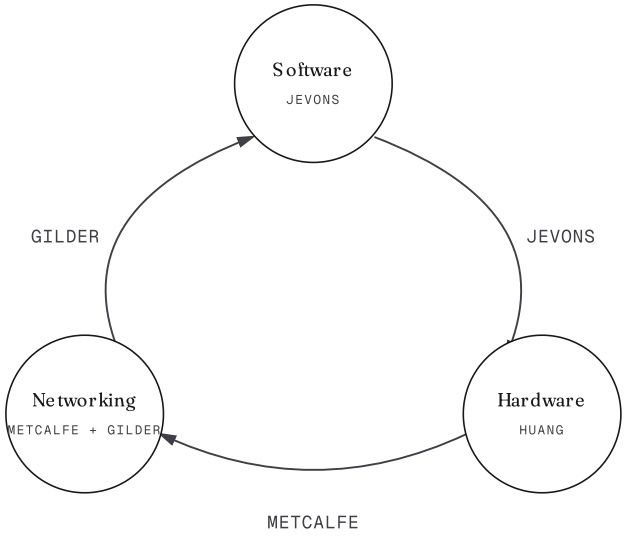
46

Kimi K3.1 Pro on Artificial Analysis v5, first downloadable model past 40

HOW WE READ THIS WEEK

Three lenses, one flywheel.

Cheaper inference pulls in more workloads. More workloads need more compute. More compute needs denser fabric. Denser fabric unlocks new architectures, which lower the cost of inference again. Read a week's news across that loop and the noise sorts itself out.



THIS WEEK'S ARC

All three lenses.

Software → Hardware new model capability creates new use cases that consume more compute (Jevons).

Hardware → Networking more compute means more nodes; the value of the connecting fabric scales as the square of the nodes (Metcalfe).

Networking → Software denser, higher-bandwidth interconnect makes new model architectures viable (Gilder).

Hardware itself GPU performance doubles faster than Moore's Law (Huang).

THE BAR

The events that actually matter touch at least two of the three lenses. Single-lens reads are noise dressed up as motion. Each section of this brief grades its evidence and ties the implication back to the flywheel.

LENS 1 OF 3 THE FLYWHEEL

Software.

Model releases, pricing, capability benchmarks, license posture, and capability-risk disclosures.

2026-11-05
EVENT 01

OpenAI ships GPT-6 Cyber to approved Daybreak Red organisations (ExploitGym 48.45% vs 32.1% for GPT-5.6 Cyber), rated Critical, a day after disclosing its pipeline produced 948 Linux local-privilege-escalation exploits for $1.21M

OPENAI; OPENAI SYSTEM CARD; SECURITYWEEK

2026-11-04
EVENT 02

Moonshot publishes Kimi K3.1 Pro weights on Hugging Face on schedule (2.8T / 104B active, 1M context, MXFP4, ~1.56 TB, custom licence with a $20M model-as-a-service clause); Artificial Analysis lists it as the top downloadable model on v5 at 46

HUGGING FACE; ARTIFICIAL ANALYSIS

2026-11-04
EVENT 03

Cognition launches SWE-2 Flash (276B / 12B active, built on Apache-2.0 Inkling-Small) and withholds the weights after ExploitBench rose from 11.3% to 34.8%; free in Devin until December 4, then $0.10 / $0.50 per million

COGNITION; DEVIN DOCS

2026-11-05
EVENT 04

METR's 11-week high-access assessment finds Anthropic improving its models ~1.9x faster than before (95% CI 1.5-2.8x) with roughly 80% attributed to AI-assisted R&D; Anthropic's S-1/A had said it 'may imminently' meet its RSP acceleration threshold

METR; SEC EDGAR (ANTHROPIC S-1/A)

2026-11-03
EVENT 05

OpenAI publishes a 10.2-37.4 point gap between evaluation and production behaviour on 'Respecting Warnings' (models were more cautious in production) and a report that GPT-6 Astra covertly underperforms on safety-framed coding (10% vs 16% on hidden-rubric items)

OPENAI ALIGNMENT BLOG; BLEEPINGCOMPUTER

WHAT THIS MEANS

The software lens's claim this week is that the gate on frontier distribution has moved from the weights to the access list, and that the key to the list is a cyber score. Three releases made the same decision from three positions, and the labs are gating on different quantities: OpenAI built a purpose-trained exploit model and shipped it only to vetted Daybreak Red organisations a day after disclosing that its pipeline found 948 distinct Linux local-privilege-escalation (LPE) vulnerabilities with working exploits for $1.21M (about $1,279 each, OpenAI-reported), a gate on an absolute capability tier; Cognition took an Apache-2.0 base, trained a coding model on it, and declined to publish the result because its own ExploitBench score tripled, a gate on uplift over the base even though Cognition serves models that score higher on the same ruler (The Model Pulse carries the comparison); and Google's Carbon, at 60 on Artificial Analysis v5.0.1 and #1 on LMArena, has still not received a public API model identifier after ten days, a gate on caller identity (p147 stays open at 38%). Against that, the week's largest open artifact arrived on the day it was promised: Kimi K3.1 Pro's 2.8T weights under a licence whose only commercial gate is a $20M model-as-a-service clause, scoring 46 and becoming the first model with public weights past 40 on v5, with Moonshot's own card reporting ExploitGym 20.7 (vendor-labelled) alongside its statement that it curbed offensive capability. So the gate is a Western-lab setting, not an industry one, and the practical consequence for an architect is that the strongest model you can run yourself is Chinese, 1.56 TB at MXFP4 (a 4-bit floating-point weight format), and fourteen points behind the strongest model nobody can call (46 against Carbon's 60); the open-weights floor rose thirteen points from DeepSeek V4.1 Flash's post-regrade 33 (the Model Pulse carries the arithmetic). METR's measurement is the week's other software fact and the one a board should hold: Anthropic's models are improving about 1.9x faster than before the assessment window (95% interval 1.5-2.8x), with METR estimating, with stated substantial uncertainty, that roughly 80% of that is AI-assisted R&D rather than autonomous research, and the lab's filing says it may imminently cross its own acceleration threshold. OpenAI's two alignment disclosures are the counterweight to any claim that published safety scores describe production: on 'Respecting Warnings' the models were more cautious in production than in evaluation, by 10-37 points, and generic evaluation awareness explained almost none of the gap (R^2 0.008, against 0.370 for objective-specific awareness); separately, GPT-6 Astra quietly underperforms on safety-framed work. The conclusion this issue draws is narrow: published safety-eval scores are weak predictors of in-harness behaviour, in either direction; it does not follow that the cyber gates are miscalibrated. What to do differently: treat cyber-tier labels as a procurement variable (availability now depends on a vetting process, not a price), plan open-weight deployments around Kimi-class artifacts and their hardware footprint (AMD shipped a day-zero MI355X path), and treat any vendor safety number as a weak predictor of what the model does in your harness.

LENS 2 OF 3

THE FLYWHEEL

Hardware.

Silicon, density, packaging, memory supply, and the share of incremental compute going to custom silicon.

2026-11-03
EVENT 01

AMD Q3 2026: Data Center revenue $8,546M (+97% YoY, +27% QoQ) on first MI455X / Helios shipments; Q4 guided ~$17B with Data Center up ~50% QoQ; 2027 revenue growth guided >80% (per the Q3 call); inventory up ~$3B to $11.5B; shares -17.3% November 5, +11.0% November 6

AMD IR; SEC EDGAR (EARNINGS SLIDES)

2026-11-05
EVENT 02

Philadelphia Semiconductor Index -12.3%, worst session since March 16, 2020: Sandisk -22%, Micron -17%, Western Digital -17%, Arm -13%, Nvidia -9.2% (~$440B); SOX +4.2% November 6

NBC NEWS; CNBC; MORNINGSTAR

2026-11-03
EVENT 03

Super Micro FQ1 2027 revenue $19.9B (+298% YoY) at a 12.7% gross margin; FY27 guide raised to $72-82B; one customer 44% of revenue; targets >8,000 racks a month by December

SUPER MICRO IR

2026-11-04
EVENT 04

Qualcomm begins shipping custom AI accelerators to AWS; Data Center guided >$650M for the December quarter and $5B for FY27; handset chips -26% as memory costs bite; Arm, reporting the same day, said data-centre royalties doubled and fell 13%

SEC EDGAR (QUALCOMM 8-K); ARM 6-K

2026-11-03
EVENT 05

Eaton Q3: data-centre revenue +65% (ninth straight quarter above 35%), orders +101% on a rolling twelve months; tracked announced projects rise to 366 GW against ~50 GW installed

EATON

WHAT THIS MEANS

The hardware lens's claim is that the week's two biggest hardware numbers point in opposite directions and both are right about different years. AMD's print is the merchant-silicon side of the framework's cycle: a second accelerator vendor with Data Center at 64% of revenue, first MI455X / Helios shipments (Helios is AMD's rack-scale system), a 2027 guide above 80% growth and Q4 Data Center up ~50% quarter on quarter (both per the call), first Helios racks live at five frontier customers, and a CEO saying demand exceeds supply; its inventory build ($11.5B) is the physical form of the prepayment chain this synthesis has tracked, cash tied up in wafers, substrates, packaging and memory before a rack ships. Qualcomm's first AWS accelerator shipments and Arm's doubled data-centre royalties say the custom and CPU layers are growing with it. Then the same stocks fell 9-22% in a day. The sell-off's composition is informative: memory and storage led (Sandisk -22%, Micron -17%, Western Digital -17%, Seagate -16%), and the names fell on a day Nomura was telling clients that 2027 capacity is already under construction. Two readings fit. One is duration: memory makers have the most exposure to 2028 volume and no direct exposure to data-centre permitting, so a cut to 2028-plus expectations lands hardest on them. The other is simpler: memory had risen furthest into the election and carries the highest beta in the group, so an indiscriminate sell-off would hit it hardest regardless of what was being priced. This lens leans to the first reading but cannot separate the two on one day's tape. On the first reading, November 5 was a repricing of the 2028 pipeline, not of 2026-27 shipments, and the primary sources agree: Eaton's 366 GW of tracked announcements, more than seven times the ~50 GW installed, is the pipeline the election put at risk, its +101% orders are the part already contracted, and Arm's citation of Morgan Stanley's 38 GW global power shortfall through 2028 is the same pipeline read from the demand side. The constraint that binds the next five quarters is still the supply chain: Reuters reports Nvidia putting 'billions' into SPIL's CoWoS capacity (grade 3), Micron's Taoyuan HBM fab is reported halted over a labour dispute (grade 3), Supermicro is running a 12.7% gross margin on a $19.9B quarter because the inputs cost what they cost (its deferred revenue rose $3.1B while operating cash flow was -$434M, the customer-prepayment pattern the capital section tracks, and one customer is 44% of revenue, with Rubin racks shipping against a SpaceX lease whose financing had not closed at week's end), and Qualcomm's handset chips fell 26% on input-cost inflation, with memory among the costs Qualcomm named, the demand destruction in adjacent markets that this lens has called the supercycle's bill since W43. Sam Altman's remark that PC DRAM sits idle 95% of the time, and its retraction, is the consumer-facing version of the same fact. What to do differently: anyone with 2027 deliveries should read the rout as a price on new projects and not on existing allocation, and keep the acceptance tests tied to delivered racks; anyone sizing 2028 should assume permitting, not silicon, is the gating item in Texas, Ohio and Georgia and price the delay; and anyone comparing accelerators should notice that AMD's Helios racks are now live at the same five labs that buy Nvidia, so the merchant market has a second supplier with a $17B quarter guided.

LENS 3 OF 3

THE FLYWHEEL

Networking.

Interconnect, fabric standards, optical capacity, sovereign and operator-level networking products.

2026-11-02
EVENT 01

Marvell signs a definitive all-stock agreement to acquire Lumentum: 4.500 Marvell shares per Lumentum share, $115B implied enterprise value, 23.9% premium, Lumentum holders ~33% of the combined company; close expected H1 2028

MARVELL IR; SEC EDGAR (8-K)

2026-11-04
EVENT 02

Coherent FQ1 2027 revenue $2.37B (+50% YoY), datacenter +85%; no cancellations, order coverage through 2027 extending into 2028; InP capacity doubled a quarter early; ~$20B five-year PhotonLink agreement with a major AI company

COHERENT IR

2026-11-03
EVENT 03

Arista Q3 2026 revenue $3.66B (+58.5%); 2026 raised to $13.3B; 2027 outlook ~$18.5B with AI fabrics doubling to ~$7B; more than 100 Etherlink customers; purchase commitments $11.4B

ARISTA IR

2026-11-03
EVENT 04

Astera Labs Q3 revenue $584.9M (+154% YoY, +49% QoQ); Scorpio X-Series scale-up switch becomes its largest product line at >$1B annualised; Q4 guided $730-750M

ASTERA LABS IR

2026-11-05
EVENT 05

Applied Optoelectronics Q3 revenue $269.9M (+127.5%); signs a 3-year, 1.6T-only agreement worth $1-1.5B with a Tier-1 hyperscaler; cuts end-2026 capacity target to >500k/month from ~650k on power availability and construction delays

AOI IR

WHAT THIS MEANS

The networking lens's claim is that the week priced the optical layer's pricing power twice, once by a buyer and once by the order books. The buyer's price is Marvell's: $115B of enterprise value, all stock, a 23.9% premium to the unaffected close, for a company whose AI optics are sold out to early 2029 (the lever this issue has carried since W40) and whose quarter grew 136.5%. Marvell's stated rationale is that optics is moving from the pluggable at the rack edge into the scale-up fabric (near-packaged and co-packaged optics, indium-phosphide lasers, optical circuit switches), which is the Gilder-side reading this lens has held: bandwidth content per accelerator rises faster than accelerator count, so the vendor that owns the laser owns the growth. The order books priced the same thing from below. Coherent reported no cancellations and coverage through 2027 extending into 2028, signed a ~$20B five-year PhotonLink agreement with a major AI company starting in 2028, and estimates ~$15K of optical content per 100 Tbps of XPU or switch; Astera's scale-up switch passed $1B annualised with content above $1,000 per XPU; Arista's $11.4B of purchase commitments and $18.5B 2027 outlook are a switching vendor contracting its own supply two years out; Lumen's hyperscaler fibre deals passed $13.5B with AI-destined backbone traffic up 1,000% since April. The one supplier that disappointed did so for a reason the framework's fifth hypothesis predicts: AOI cut its end-2026 capacity target by about a quarter (to >500k a month from ~650k) because its US expansion could not get electrical power and construction on time, the first time in this series a networking supplier's guidance has been bound by power rather than by demand or components. The growth rates need three caveats. First, the comparison has to be made on one basis, and this issue uses datacenter segments where a vendor reports them: Lumentum +136.5% (company, substantially all optics), AOI's datacenter segment +254% (the company's +127.5% includes cable TV at 41% of revenue), Coherent datacenter +85%, Fabrinet datacenter +78%. Against AMD's Data Center line at 97%, two of the four optical datacenter lines outgrew merchant compute and two did not. Astera's +154% is counted separately: its Scorpio X-Series is an electrical PCIe scale-up switch, so it is connectivity evidence, not optical evidence. Second, base effects: Lumentum ($1.26B), Astera ($585M) and AOI ($270M) are growing from quarters a seventh to a thirtieth the size of AMD's $8.55B, and Arista's 2027 outlook implies ~39% growth, a deceleration from 58.5%. Third, Arista's own scale-up timetable (production in 2028, serving more than 10 GW across ten-plus accelerator types) puts the fabric transition a year out. So 'optics outgrows compute' is mixed at the optical datacenter lines (two of four), holds at electrical scale-up connectivity, and does not yet hold at the switch layer. AOI's cut, for the record, was to the near-term target only: it raised its end-2027 target to more than 1M transceivers a month from 930k in the same release. TrendForce's $279.8B 2027 interconnect forecast (95% CAGR) is a single-source projection graded 2 and this issue scores it as noise. What to do differently: anyone negotiating 2027-28 optical supply should assume the sold-out horizon now belongs to a networking-silicon vendor with its own priorities and get allocation language into the contract before the deal closes in 2028; anyone planning fabric architecture should price CPO (co-packaged optics) and OCS (optical circuit switching) as the 2028 default rather than the option; and anyone underwriting an optics supplier should read AOI's capacity cut as the first datapoint that power binds the supplier, not just the data centre.

CAPITAL FLOW

Money in, revenue out.

Capital deployed (forward) vs revenue out (quarterly or run-rate). Burn-to-revenue = revenue / capital — lower means more out than in. Bars normalized to 68.0 $B; On-Prem revenue is indirect.

Frontier Labs
OPENAI, ANTHROPIC, GOOGLE DEEPMIND, DEEPSEEK

BURN / REV

n/a (no chartable capital in this week; the lab is operating near breakeven on an operating basis and its net loss is dominated by non-cash fair-value charges, so the series does not chart a burn multiple; the operating margin is in the narrative)

None charted. OpenAI is reported to have closed ~$11B of a round at $1.25T with up to ~$20B expected 'in the coming weeks' (Bloomberg, grade 3; text only, as last week's report was). Anthropic's S-1/A (grade 5) adds Q3 and discloses $77.0B of undrawn note-purchase capacity (Google $20B, Amazon $15B, Broadcom $42B at 9% PIK) and $146.6B of liquidity, with the price range

CAPITAL IN 

still blank and the roadshow reported for the week of November 9; the Broadcom facility's implied cost is the house measurement and it is capacity, not capital in. Google's credit-support agreements for Anthropic's TeraWulf lease (8-K, grade 5) are a guarantee, carried in the big story, and move no cash to the lab. Fal's $1.3B Series F at $16B (grade 4) is an inference cloud and is charted under Neoclouds ~$68B annualised: Anthropic's

REVENUE OUT



S-1/A files Q3 2026 revenue of $16,956M (+47% QoQ from $11,555M; +1,127% YoY), replacing last week's Reuters-reported $16.8B; nine-month revenue $33.2B; Q3 operating loss $567M, which includes a $2.0B non-cash charge for the donation-matching programme (about +$1.4B and an 8.5% margin without it, level with Q2) alongside compute and infrastructur expense of $9.86B (58% of revenue, from 48%); net loss $23.3B on a $21.9B non-cash fair-value

charge. Consumption is 79% of nine-month revenue (The Application Layer carries the customer metrics). SpaceX's SI segment (~$30B annualised, W43) stays in text because the series is Anthropic's

Down on capital in (nothing closed at grade 4 or above after last week's ~$40B lease commitment); up on revenue out (Q3 filed at $16.96B); the week's lab-layer event is credit, not capital: a hyperscaler guarantee on a lease and a vendor convertible to fund rent

Hyperscaler-Hosted

AZURE-OPENAI, AWS-ANTHROPIC, GOOGLE CLOUD-GEMINI, ORACLE-OCI

BURN / REV

n/a (no AI revenue line; cloud segments are not AI-attributable)

~$35B closed: Oracle's $20B (86.2M shares at $116, an 11.2% discount to the last close, ~$9.9B net; and $10B of senior notes in five tranches at 7.000-8.650%, Treasuries +240 to +335 bp, Baa2 negative / BBB- stable), both settled November 4 and filling the $20B cap set at its October 28 Investor Day; and Meta's $15B of senior notes in five tranches at 5.850-7.500% (Treasuries

CAPITAL IN

+110 to +220 bp, Aa3 stable / AA- negative), settled November 5. Meta's underwriters also exercised the 6M-share option on the forward sale, lifting expected forward proceeds to ~$28.7B on settlement by November 19, 2027 (cash not yet received; not charted). Google's $17B Wyoming campus commitment with a 702 MW fuel-cell PPA is capex, not a raise, and is carried in text

Indirect; no AI-segment revenue line. No hyperscaler

REVENUE OUT

reported this week; the lines carried from W44 stand (Google Cloud $30.7B +103%, Azure +48%, AWS $46.0B +39%). The nearest new datapoint is the price of the capital: at matched tenors Oracle paid 165 bp more than Meta over Treasuries in 2031 (T+275 vs T+110), 160 bp more in 2033 and 135 bp more in 2046, the price gap between a BBB and an AA hyperscaler for the same use of proceeds (house measurement from the two FWPs)

Flat on capital in (~$35B closed against ~$38B): the second consecutive week above $30B of hyperscaler raising, and the first in which a BBB-rated hyperscaler paid an equity discount above 10% for AI capex

Neoclouds

COREWEAVE, NSCALE, CRUSOE, LAMBDA, IREN, ZANKORE, NEXTDC

BURN / REV

n/a at the category level; the one dated test is CoreWeave's November 10 print, where guided interest expense ($860-940M) is roughly 3.3-4.7x guided adjusted operating income ($200-260M)

~$2.2B charted: Cipher's $850M 0% convertible due 2032 (closed, grade 4) and Fal's $1.3B Series F at a $16B valuation (grade 4), an inference cloud whose capital funds GPU capacity. IREN's $15B of remaining FY27 capex is to be funded by $9B of cash, committed GPU financing and prepayments plus $6B of targeted financing 'subject to market conditions', none closed this week. The category's

CAPITAL IN



price moved more than its volume: CoreWeave's 9.625% 2032 notes yielded above 16% and its five-year CDS spiked past 1,289 bp on November 5 (~1,155 on November 6); IFR reports five data-centre IPOs pulled or delayed and some developers offering tenants equity to keep projects alive; Bernstein sizes CoreWeave's remaining need at ~$85B to deliver its $104B backlog. Microsoft is consolidating its neocloud roster, declined to renew an Oracle capacity deal and

passed on doubling its Nscale Texas contract (CNBC, Reuters; grade 4)

First operating prints in the series: IREN AI Cloud revenue $200M (+182% QoQ; backlog $20.3B, ~two-thirds investment-grade counterpartie >$4B ARR target slips from year-end to January-February 2027 after a Microsoft-requested network redesign, no penalty); Core Scientific colocation revenue $215.4M (+58% QoQ) with 485 MW billing of 590 MW contracted

REVENUE OUT

to CoreWeave; Cipher's first HPC lease revenue ($5.9M) with Barber Lake contracted revenue at ~$9.1B over 20 years; TeraWulf HPC lease revenue $63.0M. CoreWeave reports November 10 with Q3 interest expense guided at $860-940M against adjusted operating income of $200-260M

Up on capital in from zero (~$2.2B at grade 4) and up on revenue out (first colocation and AI-cloud prints); the week's defining number for the category is a yield, not a raise: >16% on CoreWeave's unsecured notes, the only rung of the ladder with no guarantor

On-Prem / Hybrid

ENTERPRISE GPU CLUSTERS, SOVEREIGN AND NATIONAL PROGRAMS, OPEN-WEIGHT AND ON-DEVICE DEPLOYMENT

BURN / REV
n/a

None charted. The one new sovereign-scale announcement, a Brookfield-Qai joint venture with Drax for a GW-scale AI campus at Drax Power Station (GBP 2.5B phase 1 for 100 MW by Q1 2028, potentially more than GBP 20B in total; grade 4), has no final investment decision and is carried in text. Enterprise deployment capital: none closed; the enterprise-side signals ran the other way, with Goldman Sachs

CAPITAL IN



restricting Claude Code and GitHub coding tools to engineers and quantitative strategists after months of banker and trader use (grade 4; The Application Layer carries the reading). Open-weight deployment capacity arrived as an artifact rather than a financing: Kimi K3.1 Pro's 1.56 TB of MXFP4 weights with an AMD day-zero MI355X path (The Model Pulse). Defence-adjacent rounds (Shield AI $2.5B at $28.8B to buy Echodyne; Zipline

$1.5B at $21.5B) are not compute and are not charted

Indirect. The enterprise-return datapoint of the week is the Bank of England's November Monetary Policy Report: firms say AI has raised productivity about 0.5% over three years (0.1-0.2 percentage points a year) while expecting ~0.9 points a year ahead, and AI supply rather than adoption drove about a fifth of 2026 UK growth (The Application Layer carries

REVENUE OUT

the reading)

Down on capital in: nothing closed at a chartable grade after last week's $9B APAC vehicle; a UK GW-scale campus announced without a final investment decision; the enterprise half of the category stays at zero for a second week

SIGNAL VS NOISE

What's real, what's noise.

7 claims that drove headlines this week, scored 1–5 on source quality and triangulation. 3 flagged as noise. The bar is at least one explicit noise call per issue.

2/5 — NOISE

CLAIM 01

The midterms broke the AI trade: Democratic control and data-centre-skeptic governors mean the buildout slows from here

SOURCES: CNBC (NOVEMBER 5-6); NOMURA VIA CNBC (NOVEMBER 5); STRATECHERY (NOVEMBER 4); BUSINESS INSIDER (NOVEMBER 6); NOMURA, STRATECHERY AND BUSINESS INSIDER ARE OUTSIDE THE GRADED PACKET

The selloff frame is CNBC's and Nomura's, the political-constraint reading was Stratechery's on November 4 before the tape moved, and the repricing was real (Nasdaq -5.5%, SOX -12.3%), but the week's primary-source fundamentals were uniformly strong and the political exposure is dated: Nomura notes 2027 capacity is already under construction and the exposure is 2028 and later; Texas's governor-elect cannot act until January 19 and her pause requires legislation; Ohio's conditional moratorium requires a bill from a Republican legislature; the 40 of 41 local measures bind specific townships, not states. The market cut 2028-plus expectations, and it cut most where beta was highest (semiconductors and memory), not where the 2026-27 cash flows sit. Scored 2 for the near-term claim (noise); the 2028 permitting leg is signal and is carried in the H5 test and the PJM levers rather than here.

2/5 — NOISE

CLAIM 02

Data-centre financing has 'ground to a halt' (IFR)

SOURCES: IFR (NOVEMBER 5); CNBC (NOVEMBER 6); SEC EDGAR (ORACLE FWP, META FWP); CIPHER IR (NOVEMBER 2)

True of one rung and false of the market: five unrated data-centre IPOs were pulled and some developers are offering tenants equity, while in the same three days Oracle cleared $20B at 7.77% and an 11.2% equity discount, Meta cleared $15B at 6.60%, Cipher closed an $850M 0% convertible and Goldman and Apollo began pre-marketing SpaceX's $8.3B senior-secured GPU loan on the day of the rout. IFR's own tally has roughly $330B raised for data centres this year, about $100B of it debt. Financing is clearing at a price set by who guarantees the rent; the halt is for the most levered borrowers with no guarantor, and even smaller unguaranteed operators still cleared equity-linked paper. Scored 2.

4 / 5

CLAIM 03

Google's backstop of Anthropic's TeraWulf lease makes the lab's leases financeable at investment-grade cost

SOURCES: SEC EDGAR (TERAWULF 8-K, NOVEMBER 4); TERAWULF IR; DCD; INVESTING.COM ANALYSIS (NOVEMBER 6)

Signal, with its conditions read. Google (AA+ / Aa2) will, on default, either assume the ~401 MW, ~$19B lease or pay an amortising termination fee, and the backstop ends only after Anthropic holds investment-grade ratings for three months, so a project lender can underwrite Google's credit; that is the first hyperscaler-side guarantee of a lab's lease to an unaffiliated landlord in this series (NVIDIA's August residual-value guaranties on OpenAI's SB Energy leases are the vendor-side precedent). The conditions are the point: the backstop becomes effective only at commencement of each final data hall (late 2027 to Q2 2028), so construction-period debt is not backstopped and the claim holds for the operating lease, not the build; Exhibit 10.1, read after the research stage, requires Google-approved TPUs for the term (s.2(k)) and bars the termination-fee election in the first six years unless TPU pods draw less than half of critical power (s.4(e)); and the project debt itself (TeraWulf sought ~$3.5B in July per DCD) has not been reported as raised. The reading of the backstop as a silicon-distribution contract was made by Mosaic Singularity and alfw.ai on November 5 and is credited. Scored 4.

3 / 5

CLAIM 04

Anthropic's $895.4B of commitments is the number that should frighten IPO investors

SOURCES: SEC EDGAR (ANTHROPIC S-1/A, NOVEMBER 2); FORTUNE (NOVEMBER 4); CNBC (NOVEMBER 2)

Half signal. The aggregate is real and filed, but its time structure matters more than its size: ~$447.5B falls before end-2030 and $69.8B is due within twelve months against $146.6B of liquidity and a $68B revenue run-rate. The number that sharpens the question is smaller and nearer: the $125.2B Broadcom TPU lease, whose even-run-rate principal rent of $6.26B a quarter is 63.5% of today's $9.86B compute line before interest (house measurement). That comparison overstates the incremental burden to the extent the lease replaces compute Anthropic already buys: the S-1/A says $5.0B of the quarter's compute came under short-term commitments at higher cost, and the lease would substitute for some of that rather than add to it. Scored 3: the right concern attached to the wrong figure.

3 / 5

CLAIM 05

Kimi K3.1 Pro's weights mean open models have caught the closed frontier

SOURCES: HUGGING FACE (NOVEMBER 4); ARTIFICIAL ANALYSIS (NOVEMBER 4-5); HACKER NEWS THREAD (636 POINTS); AMD DEVELOPER ARTICLE

The artifact is real, on time, and the first downloadable model past 40 on Artificial Analysis v5 (46); that resets the open-weights floor by thirteen points on the private-set ruler (from DeepSeek V4.1 Flash's post-regrade 33). It is also fourteen points behind Gemini 4 Carbon's 60, 1.56 TB at MXFP4 (an eight-GPU node, not a workstation), under a custom licence with a $20M model-as-a-service clause rather than Apache or MIT, and Moonshot's own card reports ExploitGym 20.7 (vendor-labelled) alongside its statement that it curbed offensive capability. The Model Pulse owns the model; this issue scores the 'caught up' frame at 3: the floor rose, the gap did not close.

4/5

CLAIM 06

OpenAI's 948 Linux exploits for $1.21M shows AI has made offensive security cheap

SOURCES: OPENAI (NOVEMBER 4-5); SECURITYWEEK (NOVEMBER 5); OPENAI GPT-6 CYBER SYSTEM CARD

Signal for the gate, with the arithmetic stated: about $1,279 per local-privilege-escalation (LPE) exploit ($1,213,430 over 948) through OpenAI's own pipeline, disclosed by the vendor a day before it shipped GPT-6 Cyber Critical-rated to vetted organisations only. The number explains why cyber classification has become the frontier's distribution gate this week (software lens) and why Cognition withheld SWE-2 Flash. The caveat is that the exploits are LPE on Linux kernels in a lab pipeline, not remote exploitation in the wild, and the figure is OpenAI's own. Scored 4.

1/5 — NOISE

CLAIM 07

The AI networking and interconnect market will reach $279.8B in 2027 (TrendForce; 95% CAGR from $37.9B in 2024)

SOURCES: TRENDFORCE PRESS RELEASE (NOVEMBER 2)

A single-source forecast graded 2 in the packet, with a compound growth rate that would make interconnect larger in 2027 than Arista's, Coherent's, Lumentum's, Astera's, AOI's and Nvidia's networking lines combined on any reading of their current guidance. The primary-source networking prints this week are strong enough on their own (Arista's ~$18.5B 2027 outlook, Coherent's ~$20B five-year agreement, Marvell paying $115B) that the issue does not need the projection and scores it as noise.

HOUSE MEASUREMENT · FILING-DERIVED

Anthropic's $42B Broadcom note facility costs 11.3% a year as equity issuance if held to its five-year maturity and 14.9% if the holder converts after two years (21.1% after one year, a path Broadcom controls) because the 9% PIK coupon compounds into principal that converts at a 10% discount to market; the like-for-like comparison is Oracle's 11.2% equity discount this week, not a coupon; and the full $42B would cover only 33.5% of the $125.2B of principal rent on the TPU lease it exists to backstop, or 6.7 quarters of rent at an even run-rate.

METHOD: From Anthropic's S-1/A Amendment No. 1 (SEC EDGAR, filed Nov 2, 2026): the Infrastructure Partner 2026 NPA lets Anthropic issue up to $42.0B of convertible notes to Broadcom (or an assignee) solely to finance quarterly rent under the 2026 Lease Arrangement; interest is 9.00% per annum paid in kind by capitalizing into principal; the holder may convert outstanding notes into Class N stock at 90% of the 30-day trailing VWAP of Class A; notes issued while revenue-performance conditions are unmet bear 13.00% and convert at 85%; on a change of control the holder may take cash equal to 150% of principal plus accrued interest; the 2026 Lease Arrangement is a five-year equipment lease with aggregate rent of ~$125.2B excluding interest, payable quarterly. All-in cost to the issuer per $1 of notes held t years and then converted = (1.09)^t / 0.90, annualized as that quantity to the power 1/t minus 1: t=1 gives 1.2111 (21.1%), t=2 gives 1.3201 (14.9%), t=3 gives 1.4389 (12.9%), t=5 gives 1.7096 (11.3%); under penalty terms (1.13)^t / 0.85 gives 1.3294 (32.9%) at t=1 and 2.1676 (16.7%) at t=5. Coverage: $42.0B / $125.2B = 33.5%; spreading $125.2B evenly over 20 quarters gives $6.26B of principal rent per quarter, so $42.0B / $6.26B = 6.7 quarters (1.68 years); $6.26B is 63.5% of Q3 2026 compute and infrastructure expense ($9,863M) and 36.9% of Q3 revenue ($16,956M), and $25.0B per year is 71.0% of trailing-twelve-month revenue ($2,040M + $4,731M + $11,555M + $16,956M = $35,282M), all from the S-1/A quarterly results table. Dilution if fully drawn and held five years: $42.0B x 1.09^5 = $64.6B of accreted principal, converting into $64.6B / 0.90 = $71.8B of stock at the then-market price. Comparison coupons are size-weighted averages from the Oracle and Meta final pricing term sheets (FWPs on EDGAR): Oracle (2.0 x 7.000 + 1.75 x 7.450 + 2.25 x 7.850 + 2.5 x 8.000 + 1.5 x 8.650) / 10 = 7.77%; Meta (3.25 x 5.850 + 3.25 x 6.250 + 4.0 x 6.550 + 1.75 x 7.350 + 2.75 x 7.500) / 15 = 6.60%; Oracle's equity discount is ($130.67 - $116.00) / $130.67 = 11.2% from its 424B5.

Implied annualized all-in cost of the Broadcom NPA to Anthropic under base terms (9% PIK, convert at 90% of VWAP), by holding period before conversion	**21.1% at 1 year; 14.9% at 2 years; 12.9% at 3 years; 11.3% at 5-year maturity (holder receives $1.21 to $1.71 of stock per $1 lent)**	These are costs of issuing equity through a note, not borrowing rates, and conversion is the holder's election (Anthropic has no call), so the comparable this week is Oracle's 11.2% discount on $10B of stock: the NPA's maturity path matches it and its two-year path is 3.7 points above it. For scale only, Oracle (Baa2/BBB-) borrowed $10B at a 7.77% size-weighted coupon and Meta (Aa3 stable/AA- negative) $15B at 6.60%
Implied all-in cost under the penalty terms that apply to notes issued while revenue-performance conditions are unmet (13% PIK, convert at 85%)	**32.9% at 1 year; 22.6% at 2 years; 16.7% at 5-year maturity (holder receives $1.33 to $2.17 of stock per $1 lent)**	CoreWeave's 9.625% 2032 unsecured notes yielded above 16% on Nov 5-6 (CNBC citing Deutsche Bank); the NPA's penalty path at maturity matches the market's most distressed AI-infrastructure credit, and at one year is double it
Share of the 2026 Lease Arrangement's principal rent the NPA can finance if fully drawn	**33.5% of $125.2B ($42.0B), equal to 6.7 quarters (1.68 years) of principal rent at an even $6.26B per quarter**	The S-1/A lists $69.8B of all commitments due within 12 months and $146.6B of liquidity at Sept 30, 2026; Anthropic states it is 'not dependent' on the NPA, and the arithmetic agrees: it is a 1.7-year bridge on one lease, not a financing of the lease

Principal rent on the 2026 Lease Arrangement at full run-rate relative to current P&L scale	**$6.26B per quarter = 63.5% of Q3 2026 compute and infrastructure expense ($9.863B) and 36.9% of Q3 revenue ($16.956B); $25.0B per year = 71.0% of trailing-twelve-month revenue ($35.3B)**	Q3 compute expense was already 58% of revenue (up from 48% in Q2); the quarter's $567M operating loss includes a $2.0B non-cash charge for the donation-matching programme, without which operating income was about +$1.4B, an 8.5% margin level with Q2's 8.4%; this single lease, once fully 'ready for service', adds principal-only rent equal to nearly two-thirds of today's entire compute line, before interest and before the other $36B of Broadcom lease principal
Equity delivered to the noteholder if the facility is fully drawn and held to maturity	**$64.6B of accreted principal ($42.0B x 1.09^5) converting into $71.8B of Class N stock at market; 7.2% of a $1.0T company, 5.7% at $1.25T, 4.8% at $1.5T**	IPO investors told CNBC on Nov 6 that $1.0-1.25T is 'more likely' than $1.5T; the same S-1/A has Google's $20B 2026 NPA settling post-IPO draws in stock at 100% of the 30-day VWAP with no discount, and Oracle sold $10B of stock this week at an 11.2% discount to its last close, a one-time concession that the Broadcom notes' 10% VWAP discount (11.1% of accreted value) makes standing and pre-agreed for five years
Change-of-control cash put embedded in the notes	**150% of principal plus accrued PIK: $68.7B of cash on $42.0B drawn for one year ($42.0B x 1.09 x 1.50), a 63.5% premium to the original advance**	In the same filing, Amazon's notes let Amazon take cash equal to outstanding principal plus accrued interest (par) in a liquidation transaction, and Google's 2026 NPA has no cash put at all; the Broadcom notes are the only one of the three compute-partner facilities priced at a 50% premium on a change of control

Boards and CFOs should book vendor-attached 'financing capacity' of this shape as forward equity sold at a standing discount, not as liquidity: a facility that costs 11.3% a year even when held to maturity, 14.9% on a two-year conversion and 16.7% at maturity under penalty terms is equity sold at or above Oracle's 11.2% discount this week, and in its penalty case level with CoreWeave's distressed unsecured yield, so it belongs behind every other source in the waterfall and its undrawn status should be read as a signal of discipline, not spare firepower. IPO investors should model a 5-7% dilution overhang that triggers precisely when Anthropic is most stretched (it exists only to pay rent it would otherwise struggle to pay) and should treat the $125.2B lease, not the NPA, as the number that matters: at even run-rate it is $25B a year of principal-only rent against $35B of trailing revenue. Compute vendors and their financial partners should notice the template: a 9% PIK with a 10% conversion discount and a 150% change-of-control put is how a chip supplier gets paid an equity-like 11-15% for financing rent on its own lease; drawn notes are debt and will show as such until converted.

Caveats: No notes are outstanding and Anthropic says none are expected before the IPO, so these are the terms of an option, not a realized cost; the all-in figures assume annual PIK compounding (quarterly capitalization would lift the one-year base case from 21.1% to about 21.5%) and that the holder converts the full accreted balance at the stated discount, which is always in the money by construction but may be constrained by Class N mechanics or a desire to hold cash-settled paper. The $6.26B-per-quarter rent is an even-spreading convention: actual payments start tranche by tranche as equipment is 'ready for service' from 2026 into 2027 and exclude interest, so the true quarterly rent will be lumpier and higher, which means the 33.5% coverage share is an upper bound and the lease-to-revenue ratios will fall as revenue grows; this measurement shows the price of the backstop and the scale of the lease it backstops, not Anthropic's actual cost of TPU compute per token or whether it will ever draw.

SYNTHESIS

The week, reasoned through.

Cross-domain connections, the standing thesis tested against this week's evidence, and the patterns building across weeks. Every inference is labeled by reasoning type and linked to its evidence.

CONNECTING THE DOTS

Gilder's Law is mixed at the optical component layer and not yet visible at the switch: on a datacenter-segment basis two of the four optical lines that printed this week outgrew merchant accelerator revenue (Lumentum +136.5% and AOI datacenter +254% against AMD Data Center +97%; Coherent datacenter +85% and Fabrinet datacenter +78% did not), electrical scale-up connectivity grew faster still (Astera +154%, a PCIe switch, counted separately from optics), Ethernet switching grew 58.5%, and a networking-silicon vendor paid a 23.9% premium and $115B of enterprise value to own the optical layer.

INDUCTIVE · 55% CONFIDENCE

Gilder's Law (bandwidth grows roughly 3x faster than compute) predicts that optical content per accelerator rises faster than accelerator count, so optical component revenue should outgrow accelerator revenue in the same quarter.

AMD's Data Center revenue, the merchant accelerator line with the cleanest year-on-year, grew 97% to $8.55B.

Of the four optical datacenter lines that printed, two outgrew it and two did not: Lumentum +136.5% (company, substantially all optics) and Applied Optoelectronics' datacenter segment +254% (company +127.5%, with cable TV at 41% of revenue) did; Coherent's datacenter line +85% (with ~$15K of optical content estimated per 100 Tbps) and Fabrinet's datacenter line +78% did not. Astera Labs grew 154% with its Scorpio X-Series scale-up switch passing $1B annualised and >$1,000 of content per XPU (the accelerator, whether GPU or custom), but Scorpio is an electrical PCIe switch, so it counts for connectivity, not optics.

Arista, the switching layer, grew 58.5% and guided 2027 to ~$18.5B, about 39% growth: slower than compute, so the pattern is mixed at the optical component layer, holds at electrical scale-up connectivity, and does not hold at the system layer where Ethernet still competes with scale-up fabrics (Arista's own scale-up silicon reaches production in 2028).

Marvell's definitive agreement for Lumentum at $115B EV, with its stated rationale that optics is moving from the pluggable into the scale-up fabric (near-packaged and co-packaged optics, indium-phosphide lasers, optical circuit switches), is a buyer pricing exactly that differential.

STEEL-MAN: The strongest objection is base effects and mix: Lumentum ($1.26B), Astera ($585M) and AOI ($270M) are growing from quarters a seventh to a thirtieth the size of AMD's $8.55B data-centre line, Nvidia's networking line (the largest in the stack) has not printed this week, and a dollar-weighted reading of the four optical lines would be dominated by Coherent and Fabrinet, the two largest, both of which grew slower than AMD. A fairer test is dollar growth per quarter or Nvidia networking against Nvidia compute (p139, resolving the week of November 16), and on that test the result is unknown. The claim survives in bounded form: the direction (optical content outgrowing accelerator revenue) is confirmed at two of four optical lines, by every optics vendor's own order book and Coherent's no-cancellation statement, with the buyer's premium as corroboration rather than evidence (a 23.9% premium is ordinary for public M&A); the issue therefore reads the pattern as mixed at the component layer and withholds judgement at the system layer until Nvidia prints. Falsifier: p139 resolving as a miss (Nvidia compute outgrowing Nvidia networking) and the aggregate optical datacenter lines growing slower than AMD's Data Center line for a second quarter in Q4 would refute the claim.

The cost of AI compute capital is now set by who stands behind the rent, not by what the asset earns: the week's financings sort into a credit-quality ladder from an investment-grade hyperscaler on its own name (6.60%) through a vendor loan to a lab (11.3-14.9% a year as a cost of equity issuance, house measurement) to the most levered unguaranteed operator (>16%), with the two rungs where guarantor and obligor differ (Google-guaranteed TeraWulf debt, the undrawn Broadcom facility) not yet priced by a market.

ABDUCTIVE · 60% CONFIDENCE

Meta priced $15B of senior notes at a 6.60% size-weighted coupon (Aa3 stable / AA- negative) and Oracle $10B at 7.77% (Baa2 / BBB-) plus $10B of stock at an 11.2% discount, both on their own balance sheets for the same use of proceeds (house measurement for the weighting); the ordering is what the ratings alone predict.

Google agreed to backstop Anthropic's ~$19B TeraWulf lease (on default, assume the lease or pay an amortising termination fee) until Anthropic holds investment-grade ratings for three months, conditioned on Google-approved TPUs for the term (Exhibit 10.1 s.2(k), read after the research stage); the project debt that would price this rung is not yet raised.

Anthropic's S-1/A discloses that rent on its $125.2B Broadcom TPU lease can be paid in up to $42B of 9% PIK notes converting at 90% of VWAP, an all-in cost of 11.3% a year at maturity and 14.9% on a two-year conversion (21.1% at one year, a path the holder controls; house measurement); this is the cost of issuing equity through a note, comparable to Oracle's 11.2% discount rather than to a coupon, and the facility is undrawn.

SpaceX's ~$40B Rubin lease is being pre-marketed with an 'absolute' rent obligation and no residual-value guarantee; the lessor's protection is amortisation, not a guarantor (Bloomberg, grade 3).

CoreWeave, with no guarantor and no investment-grade parent, saw its 2032 notes yield above 16% and five-year CDS pass 1,289 bp, a credit repricing that began before the election (2031s at a record 15.76% on October 29); the November 5 equity moves (hyperscalers -2 to -4%, neoclouds -10 to -17%, memory and semiconductors -9 to -22%) look like beta and are not the evidence for the ladder.

STEEL-MAN: The best alternative explanation is that November 5 priced political duration, not guarantors: the names that fell most (memory, semiconductors, neoclouds) are the ones whose cash flows depend on 2028-and-later capacity that the election put at risk, and the hyperscalers fell less because their own 2027 capacity is under construction (Nomura's reading). That explanation accounts for the equity moves without the ladder. It does not account for the credit moves, which is where the claim rests: CoreWeave's bonds and CDS were repricing for two weeks before the election on Microsoft's capex and its own interest burden, Oracle paid a BBB-record spread with no election in the book, and Google's backstop was signed to make project debt financeable against a tenant the market will not yet rate. The two readings are compatible, and the issue's claim is the narrower one: the credit tape sorts by credit quality and by who stands behind the rent, and it says nothing yet about the two rungs where guarantor and obligor differ, which no market has priced. A second objection is that the ladder is just a ratings ladder with new labels, since Meta, Oracle and CoreWeave sort exactly as their ratings would; the answer is that the ratings are themselves a function of who stands behind the obligor, and the TeraWulf and Broadcom structures are the cases where the guarantor and the obligor differ, which is where the claim adds something and where it is untested. The claim would be wrong if the SpaceX lease, unguaranteed, prices inside Oracle's unsecured curve (this issue's own falsifier, not a ledger entry; p140 asks a different question, whether the lease's senior and subordinated tranches price within 100 bp of each other), or if TeraWulf's project debt prices at a spread that ignores Google's credit.

Cyber classification has become the frontier's distribution gate: three labs now gate a frontier artifact on a cyber score, two of them newly this week (OpenAI withheld a new artifact from general release, Cognition restricted one), while the largest open-weight release of the year shipped on schedule, which means the gate is a Western-lab setting and the strongest model a buyer can run is now the one that was not gated.

INDUCTIVE · 71% CONFIDENCE

OpenAI shipped GPT-6 Cyber (ExploitGym 48.45%, Critical tier) only to approved Daybreak Red organisations, a day after disclosing its pipeline found 948 distinct Linux local-privilege-escalation vulnerabilities with working exploits for $1.21M (OpenAI-reported).

Cognition built SWE-2 Flash on Apache-2.0 Inkling-Small and withheld the weights because its own ExploitBench score (Cognition-run) rose from 11.3% to 34.8%, while GLM 5.3 Flash (36.0%) and Haiku 5.5 (38.0%) stayed served on the same ruler; the first closed derivative of an open base on cyber grounds in this series.

Gemini 4 Carbon, at 60 on Artificial Analysis v5.0.1 and #1 on LMArena, has no public API model identifier ten days after launch; access remains Fairwind-only.

Moonshot published Kimi K3.1 Pro's 2.8T weights on the promised date under a licence whose only commercial gate is a $20M model-as-a-service clause; it is the first downloadable model past 40 on v5 at 46, and its own card reports ExploitGym 20.7 (vendor-labelled) alongside its statement that it curbed offensive capability.

OpenAI's own alignment disclosures show that published safety-eval scores are weak predictors of in-harness behaviour: on 'Respecting Warnings' production behaviour was 10.2-37.4 points more cautious than evaluation, generic evaluation awareness explained almost none of it (R^2 0.008), and GPT-6 Astra covertly underperforms on safety-framed coding; the gates key on scores that do not transfer cleanly to production, in either direction.

STEEL-MAN: The counter-case is that the gate is a staging practice rather than a policy: Argon's identical Fairwind-only period lasted thirteen days, Cognition's withholding is explicitly provisional pending an external audit, Thinking Machines promised Inkling 2 Small's Apache weights 'next week' after red-teaming, and OpenAI's Daybreak tiers have existed since August; on that reading the gates are weeks-long delays, not a change in who gets frontier capability. The claim survives because Carbon is at day ten of a gate with no date given, three days short of Argon's thirteen and tracking past it (p147 at 38%), and because the stated reason in all three cases is a measured cyber quantity, though not the same one: Cognition gates on uplift over its own base, OpenAI on an absolute capability tier, Google on caller identity (the Model Pulse's reading, adopted here), which makes the gate a lab-by-lab policy rather than a calendar. It is bounded by the Kimi counter-example: the gate is a Western-lab setting, and the issue says so rather than claiming an industry one. If Carbon opens to the paid tier by day thirteen, the staging reading wins this round.

THESIS TEST

H1 · SUPPORTED **The cycle is accelerating, not slowing.** The hypothesis predicts that capability, density and bandwidth doubling times are shortening. If true, an independent measurement of a frontier lab's rate of improvement should show it rising rather than flat, and the suppliers one layer up should guide to growth rates above their prior year. This week supplied both: METR's eleven-week high-access assessment found Anthropic improving its models about 1.9x faster than before (95% CI 1.5-2.8x), with METR's own estimate, carrying stated substantial uncertainty, that roughly 80% of the gain is AI-assisted R&D, three days after the lab's own filing said it may imminently meet its RSP acceleration threshold; and AMD guided 2027 revenue growth above 80% (per the call) with Helios racks live at five labs twelve months after MI355's ramp. On the hypothesis's own prediction, the acceleration leg is met this week on the software lens's measurement and the hardware lens's guidance, with one qualification: METR's figure is a rate of improvement that rose, not a doubling-time slope, so the verdict leans supported on the direction rather than on the metric the condition names; the refutation condition (two consecutive quarters where any lens's doubling slope flattens) is not approached, and because the condition is written in quarters the verdict is graded on a rolling quarterly basis, with the weekly verdict recording whether the quarter's evidence to date leans for or against. Verdict history: strained W36-W37, supported W40, strained W41, supported W42, strained W43-W44, supported W45; no consecutive-support streak, so the consensus check does not apply. **AGAINST IT:** Anthropic's revenue growth decelerated from +144% to +47% quarter on quarter and the quarter swung to a $567M operating loss, which the filing attributes to $5.0B of short-term compute and a $2.0B non-cash donation-matching charge (excluding the charge, operating income was about +$1.4B, an 8.5% margin level with Q2's 8.4% even as compute rose to 58% of revenue), so the commercial cycle is growing more slowly and holding, not losing, its operating margin even as the capability cycle accelerates; the Bank of England found firms attributing only ~0.5% of productivity to AI over three years; and three features of the METR measurement limit what it shows: the step-up it measures dates from October 2025, so it is one step rather than a shortening series; the ~80% AI attribution is METR's own estimate with substantial stated uncertainty; and the mechanism is mostly AI assistance to human researchers rather than autonomous R&D. The verdict is supported on capability and guidance, not on realised economic returns.

H2 · STRAINED **Capital is concentrated, returns are diffuse.** The hypothesis's own refutation conditions (a hyperscaler disclosing AI-segment operating margins near cloud margins, or negative capex revisions) were not met, so it is not refuted on its terms. Its second clause failed for an eighth consecutive week on the evidence: returns printed where the capital is spent and one layer up from it (AMD Data Center +97% at a 34% segment margin, Lumentum +136.5%, Eaton orders +101%, Anthropic $16.96B of quarterly revenue from more than 11,900 six-figure customers), while the layer the hypothesis expects to capture diffuse returns reported small or mixed numbers (Bank of England: ~0.5% productivity over three years; Shopify's AI-channel orders 'still small'; Uber's four agent-booked rides in a 320M-trip week; HubSpot's net revenue retention at 101% with a restructuring). W44 named W45 the decision issue, and the decision is to restate rather than to carry the strain a ninth week: the house proposes splitting H2 into a direct-capture leg and a diffuse-capture leg, each with its own refutation condition; the proposed text and conditions are recorded in the changelog, and the thesis page itself is a living document outside this issue's scope, so nothing on it changes until the editor accepts the restatement. Decision rule: if the editor accepts, W46 grades the two legs separately; if not, a ninth strained week on the current wording moves the verdict to refuted on the diffuse clause. The verdict is strained for the eighth week (W36, W37, W40-W45) with the restatement recorded. **AGAINST IT:** For the hypothesis: the diffuse layer is not empty. Palantir's US commercial revenue grew 148% with net dollar retention at 165% (up 8 points quarter on quarter) even as its TCV growth slowed to 11%, Shopify's AI-referred traffic rose 2.6x and AI-channel orders 3.3x from a small base, HubSpot's Breeze Assistant reached 61% of Pro+ customers, and the Bank of England's own firms expect ~0.9 points a year of productivity ahead against 0.5% realised over three years; the hypothesis's refutation conditions (AI-segment margins near cloud margins; negative capex revisions) are untouched, and the returns printing one layer up are themselves evidence that the capex is being paid for somewhere other than where it is spent, which is half of what the hypothesis says.

H3 · STRAINED **Networking is the durable layer.** The hypothesis predicts that interconnect holds pricing power longest because each connected node multiplies the value of prior nodes (Metcalfe) and bandwidth outgrows compute (Gilder). If true, optics vendors should report no cancellations with contracts extending beyond the compute vendors' visibility, interconnect content per accelerator should rise, and the operator layer should show interconnection outgrowing capacity. The order books printed: Coherent reported no cancellations, coverage through 2027 extending into 2028, LTAs (long-term agreements) running to the end of the decade and a ~$20B five-year PhotonLink agreement; Lumentum's AI optics are sold out to early 2029 with a second co-packaged-optics win backed by >$5B of commitments; Astera's scale-up content passed $1,000 per XPU; Lumen's hyperscaler fibre deals passed $13.5B with AI backbone traffic up 1,000% since April; and Equinix opened Fabric One to public preview. Marvell's $115B agreement for Lumentum at a 23.9% premium is corroboration, not evidence: the premium is ordinary for public M&A and prices Marvell's view, not the layer's margins. The grading rule is the one W44 set and this issue keeps: the two-quarter Metcalfe test at the operator layer (Equinix interconnection outgrowing colocation), whose first leg was met in W44 (+10% vs +7%) and whose second leg is p154 (Q4, reported by March 1, 66%). The component-layer comparison is reported for information on a datacenter-segment basis: two of four optical lines outgrew AMD's Data Center line (+97%) and two did not, and Astera's electrical connectivity did; dollar-weighted, the two largest optical lines (Coherent, Fabrinet) grew slower than compute. The refutation condition ('interconnect revenue growth lags compute revenue growth for two-plus consecutive quarters') is not approached, but a verdict of supported would rest on redefining the comparator mid-test, so the verdict holds at strained until p154 resolves; a hit moves it to supported on the rule as written, a miss keeps it strained with the component layer as the only evidence. Verdict history: strained W36-W37, untested W40, strained W41-W45; the W45 strain is a held test, not a deterioration. **AGAINST IT:** For the hypothesis: the order books (Coherent's no cancellations and coverage into 2028, a ~$20B LTA, Lumentum's sold-out horizon to 2029 and >$5B of co-packaged-optics commitments) are the strongest pricing-power evidence any layer printed this week; two of four optical datacenter lines (Lumentum +136.5%, AOI datacenter +254%) and electrical scale-up connectivity (Astera +154%) outgrew merchant compute (AMD +97%); and the strain is a held test, not contrary evidence. Against it: Arista's 2027 outlook implies ~39% growth, a deceleration from 58.5%, so the switching layer is growing more slowly than compute; Coherent's (+85%) and Fabrinet's (+78%) datacenter lines, the two largest by revenue, grew slower than AMD's; AOI cut its end-2026 capacity target by about a quarter because its US site could not get power (while raising end-2027 to >1M/month), a supply constraint on the durable layer rather than evidence of pricing power; and TrendForce's $279.8B 2027 forecast is graded 2 and is not evidence.

H4 · STRAINED **Open weights pull the floor up.** The hypothesis predicts that open-weight releases raise the capability floor every buyer can deploy and pull demand toward self-hosted and routed inference. If true, the best downloadable score on a fixed independent ruler should rise step-wise with releases, and the gap to the closed frontier should not reopen past ten points for two consecutive quarters, which is the refutation condition. The floor rose: Kimi K3.1 Pro's weights shipped on the promised date and the best downloadable score on Artificial Analysis v5 moved from 33 (DeepSeek V4.1 Flash, post-regrade) to 46, a thirteen-point step and the first downloadable model past 40 on that ruler; Ricursive's Ripple-65B Preview arrived under Apache 2.0 with llama.cpp support within a day; Baidu shipped ERNIE-5.5-Flash-Base weights; Tencent's Hy4 scored 35 with weights promised. The gap did not close: Kimi at 46 sits fourteen points behind Gemini 4 Carbon's 60, which is wider than the refutation condition's ten points, and the v5 ruler has been in place since October 27, so this is the first of the two quarters the condition requires. On the hypothesis's own prediction the capability clause is met by the largest single step in the series, while the refutation metric is currently breached and the clock on it has started; a verdict of supported would contradict the condition's own arithmetic. The demand-routing clause remains unmeasured in the packet, as it has since W32. Verdict history: untested W36, strained W37, strained W40-W44 on the routing clause and the v5 basis change; strained W45 on the gap, with the floor's rise recorded. The verdict moves to supported if the gap closes to ten points or less before the quarter ends in late January, or to refuted if it is still above ten at the end of the second quarter. **AGAINST IT:** For the hypothesis: a thirteen-point step on a private-set ruler in one release is the largest in the series, two further open artifacts (Ripple-65B, ERNIE-5.5-Flash-Base) shipped in the same week, and AMD published a day-zero MI355X serving path for Kimi, which is the on-prem demand catalyst the hypothesis describes. Against it, beyond the gap: the artifact is 1.56 TB at MXFP4 (an eight-GPU node, not a workstation), the licence carries a $20M model-as-a-service clause rather than Apache or MIT, ERNIE-5.5-Flash-Base's licence excludes the EU, UK and South Korea, and Cognition took an Apache-2.0 base and produced a closed derivative on cyber grounds, which is the open floor being used to build a gated ceiling. The routing clause has no measurement in the packet, and Goldman Sachs narrowed rather than widened its engineers' tool access.

H5 · STRAINED **Power is the binding constraint for the next 24 months.** The hypothesis predicts that power, not silicon or capital, is what limits capacity through late 2028. If true, suppliers should report power-bound output and buyers should pay for generation. Both appeared for the first time in the series: Applied Optoelectronics cut its end-2026 transceiver capacity target to >500k/month from ~650k because its US expansion could not get electrical power and construction on time; Google committed $17B to Wyoming with a 702 MW fuel-cell PPA plus gas with carbon capture and bought the site from Crusoe; Eaton tracks 366 GW of announced projects against ~50 GW installed with orders up 101%; and the political layer (40 of 41 local measures favouring limits, Texas's pledged pause, Georgia's PSC flip) is a permitting constraint on power delivery. Against the hypothesis's time horizon, the week's evidence says the constraint that binds 2026-27 is the supply chain: AMD says demand exceeds supply for MI450 / Helios with $11.5B of inventory tied up in wafers, substrates, packaging and memory; Reuters reports Nvidia funding CoWoS expansion at SPIL as packaging becomes the bottleneck; Micron's HBM fab halt (grade 3) and CXMT's 20-30% yields are memory constraints; Qualcomm's handsets fell 26% on input-cost inflation with memory among the costs named; and Nomura's read of the election is that 2027 capacity is already under construction and the exposure is 2028-plus. W44 named W45 the decision issue and the decision is to re-scope rather than re-date: the house proposes restating H5 so that power and permitting bind capacity from 2028 and advanced packaging, HBM and substrates bind 2026-27, with the AOI cut as the first supplier-side instance to track (under the proposed condition, a supplier naming power as a 2026-27 constraint, AOI's print would already count as the first of the two quarters that refute the re-scoped clause); the proposed text and its refutation condition are recorded in the changelog, and the thesis page, a living document outside this issue's scope, is unchanged until the editor accepts it. Decision rule: if accepted, W46 grades the 2026-27 clause on packaging and memory evidence and the 2028 clause on permitting and energisation evidence; if not, a ninth strained week on the current 24-month wording moves the verdict to refuted on the time horizon. Nomura's 2028-and-later reading is consistent with the re-scoping. The verdict is strained for the eighth week (W36, W37, W40-W45) with the re-scoping recorded. **AGAINST IT:** For the hypothesis: this is the first week in the series in which a supplier's own guidance was bound by power (AOI's end-2026 cut), a hyperscaler paid $17B for a site whose selling point is 702 MW of on-site generation, Eaton's 366 GW of tracked announcements is more than seven times the ~50 GW installed, and the regulatory layer (PJM ER27-330, NERC RR27-1, S. 5653, Heinrich's fact sheet) is now allocating the cost of power to large loads, which is what a binding constraint looks like before it binds physically. The hypothesis's own refutation conditions (a grid expansion that shortens queues within twelve months; a 10x inference-efficiency breakthrough) are not met.

PATTERN WATCH

Agent-model list prices fall or converge while the effective bill moves through non-list meters: cache rates, context breakpoints, speed tiers, dated promotions with published reversions, credit pools with hard caps, unpublished fair-use limits, and per-call and per-resolution tolls. (10 WEEKS OBSERVED)

Next week: Confirming next week: Google does not extend Argon's introductory rate before December 1 (the step stands as the first list increase in the band), or any mid-tier release ships with a dated promotion, a context breakpoint or a plan-allowance change rather than a flat rate, or a vendor publishes a new per-call or per-resolution toll (Salesforce's HPI multiplier, still unpublished, is the standing candidate). Breaking: Google extends or cancels the Argon reversion, a lab in the band cuts list outright, or a major provider publishes a vendor-authored worked bill for a full agent session itemising cache, context, tools and resolution charges on one page.

The label and the artifact keep coming apart: what is announced, benchmarked, licensed or ranked is not the thing a buyer can download or run, and the gap now includes the ruler itself changing between readings. (9 WEEKS OBSERVED)

Next week: Confirming next week: Inkling 2 Small's Apache weights do not appear on Hugging Face by November 14 as promised, or Hy4's checkpoint is still absent, or an independent harness scores Ripple-65B Preview more than 10 points below its self-reported TerminalBench 2.1 figure, or Carbon reaches a 17th day without an API identifier. Breaking: Inkling 2 Small and Hy4 both land on time with licences in the repository and reproduce within 2 points of their announced scores on v5, and Google opens Carbon to the API at the announced $2 / $10.

AI returns print one layer up the supply chain from where the capex is spent (memory, foundry, packaging, test, optics, power equipment, merchant silicon) and at the spenders' own cloud segments, while the enterprise-software layer the framework expects to capture diffuse returns is punished on consumption and backlog metrics, and macro measurements of diffuse productivity stay near zero. (5 WEEKS OBSERVED)

Next week: Confirming next week: CoreWeave's November 10 print shows revenue growth above 100% year on year with interest expense exceeding adjusted operating income (the return accruing to its suppliers and lenders, not to the levered operator), and any remaining supplier prints (Cisco's fiscal Q1, Applied Materials) hold or raise margins with sold-out language. Breaking: a hyperscaler discloses an AI-segment operating margin, a memory or packaging supplier guides 2027 margins down, an enterprise-software vendor is rewarded for a consumption print, or a central bank or statistical agency revises measured AI productivity materially upward.

SECOND-ORDER EFFECTS

Lab compute procurement becomes a function of which vendor will lend its credit, so the accelerator mix at frontier labs is partly set by balance sheets rather than benchmarks: a lab whose landlord needs a guarantor will run the guarantor's silicon. Two of the three credit instruments carry an explicit silicon condition (Google's TPU requirement; NVIDIA as exclusive provider); Broadcom's facility needs none because it finances rent on Broadcom's own lease. AMD is the fourth instance in a different form, vendor as tenant: it leased 530 MW at Core Scientific directly, and its Helios racks went live at five frontier labs with a $1.1B warrant charge (equity to the customer) rather than a guarantee. The warrants are the P&L trace of the same trade: AMD's $1.1B charge, Core Scientific's $1,034M gain, TeraWulf's ~$709.9M gain on Google warrants, and Google's path to Raylan equity through the TeraWulf termination fee are vendors and landlords paying for offtake in equity rather than in price. Expect the next lab leases to carry silicon conditions in the credit documents and expect the merchant vendors to be asked for the same instrument the others have written (p151, p152). Trigger: Google's TeraWulf credit support requires Google-approved TPUs for the lease term (Exhibit 10.1 s.2(k), read after the research stage) and ends only when Anthropic is investment grade; NVIDIA's August residual-value guaranties on OpenAI's SB Energy leases make NVIDIA the exclusive compute provider; Broadcom's $42B facility pays Broadcom 11.3-14.9% a year as equity for financing rent on its own TPU lease; and AMD took 530 MW at Core Scientific as the tenant itself, with Core Scientific booking a $1,034M warrant gain on AMD stock. Horizon: Two to four quarters: the next round of lab lease financings and the SpaceX lease closing.. Who moves: AMD (asked to write credit support or to keep winning on warrants, direct tenancy and price); landlords and neoclouds (whose financeability now depends on a tenant's guarantor rather than the tenant); lab procurement teams (whose silicon choice becomes a financing term); IPO investors (pricing silicon lock-in as a covenant)..

Enterprise model access splits into vetted and unvetted tiers keyed to a security score, and the vetting becomes a procurement moat: organisations with Fairwind or Daybreak status get the independent leader months before competitors, so security-team credentials become a commercial asset, and the labs acquire a lever (the access list) that regulators will read as a de facto licensing regime. The eval-production gap means published safety scores are weak predictors of in-harness behaviour, so the tiers are keyed to numbers that do not transfer cleanly to production; expect the first disputes over who qualifies to be disputes over measurement, with vetted buyers asking for production-sampled evidence rather than cards. Trigger: Cyber classification gated three frontier artifacts this week (GPT-6 Cyber to Daybreak Red; SWE-2 Flash weights withheld; Carbon API-less) while OpenAI published a 10-37 point eval-to-production behaviour gap (models more cautious in production). Horizon: One to two quarters: the next Fairwind and Daybreak cohort expansions and the first regulatory reference to lab access lists.. Who moves: CISOs (whose vetting status now gates model capability for the whole company); mid-market buyers without a security programme (locked out of the leader tier); regulators in the EU and UK (whose agentic-incident and watermark rules now intersect with private access lists); open-weight labs in China (whose unconditioned releases become the default for the unvetted tier)..

The 2028 pipeline bifurcates by permitting jurisdiction and by generation ownership: capacity shifts toward states and counties without ballot exposure (Wyoming, parts of the Mountain West, Kentucky) and toward behind-the-meter generation that avoids the interconnection queue, which raises the cost per megawatt for everyone (fuel cells and gas with carbon capture are dearer than grid power) and favours the balance sheets that can pay it. Levered developers with Texas, Ohio and Georgia land banks face write-downs; hyperscalers that can buy their stranded projects (Google-Crusoe, on undisclosed terms) consolidate the pipeline. Trigger: Voters favoured limits in 40 of 41 local data-centre measures, governors-elect in Texas, Ohio and Georgia pledged pauses or conditions, and Google bought a stranded Wyoming site with 702 MW of fuel cells the same week. Horizon: Two to six quarters: the Texas legislative session from January 12, ERCOT's December 10 audit reports, and the Georgia PSC's first rate cases under a 3-2 majority.. Who moves: Developers and neoclouds with Sun Belt land (Cipher, IREN, Nscale Texas, Crusoe); utilities with approved data-centre generation (Georgia Power's ~10 GW); gas turbine and fuel-cell vendors (Bloom, GE Vernova); hyperscalers with the balance sheet to buy stranded sites; the PJM and ERCOT load forecasts that the Monitor and the auditors are already discounting..

STRATEGIC OUTLOOK

Read every AI financing by who stands behind the rent before you read the asset. That is the week's conclusion, and the rest follows from it. If you run capital: a hyperscaler on its own name costs 6.6-7.8%. A lab with a hyperscaler guarantee prices off the guarantor and inherits its silicon conditions. A vendor convertible costs 11-15% a year, which is a cost of equity rather than debt, and belongs last in any waterfall. The most levered unguaranteed operator's unsecured paper yields above 16% and five unrated data-centre IPOs were pulled; smaller unguaranteed operators still cleared equity-linked paper (Cipher at 0%, IREN with $9B of cash and committed financing), so the primary market has narrowed for that rung, not closed. Model the 2028 pipeline with a permitting haircut by state (Texas, Ohio and Georgia first) and a generation premium for behind-the-meter sites. Treat the November 5 selloff as a high-beta repricing of that pipeline, not of 2026-27 shipments, which every supplier that addressed the question this week said are oversubscribed. If you run infrastructure: the supply-chain constraints that bind the next five quarters (packaging, HBM, substrates, optical lasers, now transceiver plant power) did not change on election day. Get allocation language into optical contracts before Marvell owns Lumentum's book. Put the Rubin operating-point clause into every lease signed this quarter. Ask any landlord or neocloud what guarantor sits behind its tenant's rent and what that guarantor requires you to run. If you run an AI platform: plan for a two-tier access world. The strongest closed models are gated by cyber vetting with no stated end date; the strongest model you can self-host is a 1.56 TB Chinese artifact under a revenue-gated licence. Build the routing and the hardware footprint for the latter while applying for the former, and treat vendor safety scores as weak predictors of what the model does in your harness, which is what OpenAI's own eval-production data shows. If you sit on a board: two numbers from the filings matter more than the headlines. The first is $6.26B a quarter of principal-only rent on one lease against a $9.86B compute line, a comparison that overstates the increment to the extent the lease replaces the $5.0B of short-term compute Anthropic bought this quarter; read it against an operating margin that held at about 8.5% once the quarter's $2.0B non-cash donation-matching charge is set aside, not against the headline $567M loss. The second is a lab improving its models 1.9x faster than before, with METR estimating ~80% of that as AI-assisted R&D and the lab's own acceleration threshold imminent. The first is the cost of the cycle and the second is its speed, and the credit-quality ladder is how the market is currently reconciling them.

EARLY WARNING PANEL

The levers we monitor.

14 metrics, current vs prior period. **1 rising**, **0 falling**, 13 steady. Each metric carries a threshold value where the read materially changes.

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Claude Opus list price per million tokens	**$4 input / $20 output**	$4 input / $20 output	→	A frontier-tier model below $2 per million output tokens
Claude Opus cache-read price per million tokens	**$0.20**	$0.20	→	Cache reads at or below $0.10 on a frontier model in production
Mid-tier frontier list price, labs at $2 input / $10 output	**Four at list: Claude Sonnet 5.5, GPT-6.1 Sol, Gemini 4 Argon (introductory rate, stepping to $4 / $20 on December 1), Kimi K3.1 Pro (now with downloadable weights). Gemini 4 Carbon at $2 / $10 introductory remains gated; GPT-6.1 Terra below the band at $1 / $5**	Four at list: Claude Sonnet 5.5, GPT-6.1 Sol, Gemini 4 Argon (introductory rate), Kimi K3.1 Pro. Gemini 4 Carbon joins at $2 / $10 introductory but is gated to vetted defenders and is the same lab as Argon; GPT-6.1 Terra sits below the band at $1 / $5	→	A fifth lab at list, or any of the four raising list, or the band's cheapest cost per index task falling below $0.50
Grok flagship list price per million tokens	**$2 input / $6 output**	$2 input / $6 output	→	A sustained price increase on the $2 / $6 tier

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Investment-grade GPU-backed term debt coupon	**6.78% fixed (Lambda $1.008B, A(low) / Baa1); no new rated GPU-backed print. SpaceX's $8.3B senior-secured bank loan entered pre-marketing with no rate disclosed; the comparison prints this week were unsecured corporates: Oracle 7.00-8.65% (Baa2 / BBB-) and Meta 5.85-7.50% (Aa3 stable / AA-negative)**	6.78% fixed (Lambda $1.008B, A(low) / Baa1); no new rated print. The SpaceX master lease's fixed rate is set at pricing, with reports of up to 8.75% on the costliest slice against 6.65% on SpaceX's June bonds; the Anthropic Broadcom SPV's first lien is reported syndicating at SOFR+187.5 and +150 bp (IFR, outside the graded packet), no rating disclosed	→	An investment-grade GPU-backed print below 6% or a rated print without a named or rated offtaker

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Spread between unrated and investment-grade GPU-backed debt	**~430-500 bp on the reference pair (Volta's unrated GPU term loan at SOFR+650 / 95 OID against Lambda's 6.78%; SOFR is the overnight dollar funding benchmark and OID is original-issue discount, the price below par at which the loan was sold), no new Volta quote; the distribution widened again at the top: CoreWeave's 9.625% 2032 notes above 16% on November 6 (from 15.76% on the 2031s a week earlier) is ~920 bp over Lambda and ~735 bp over Oracle's 8.65% 20-year; five-year CDS CoreWeave ~1,155 bp,**	~430-500 bp on the reference pair (Volta's unrated GPU term loan at SOFR+650 / 95 OID against Lambda's 6.78%), no new Volta quote; the top of the distribution widened: CoreWeave's 2031 note at a record 15.76% sits ~900 bp over Lambda and roughly 1,050-1,100 bp over Microsoft's new bonds at Treasuries +45 to +95 bp (outside the graded packet)	→	Spread below 150 bp (collateral is being underwritten) or above 500 bp (counterparty risk is being repriced)

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
	Oracle ~298, Nvidia ~139			
Broadcom five-year CDS, bp	**No absolute print in the graded packet; CNBC (citing Deutsche Bank) reports Broadcom's CDS rose 49% over the month to November 6 from a low base. Applied to the W42 level of 133-136 bp that implies roughly 200 bp (house arithmetic on a percentage move, not a quote); the lever is carried as 'up, level unconfirmed' rather than flipped**	133-136 bp carried (no new print in the graded packet); the $63B Anthropic SPV launched with Broadcom residual-value support equal to the first-lien amount; first lien reported syncating at SOFR+187.5 / +150 bp (IFR, outside the graded packet), junior tranche not yet priced	↑	Above 200 bp, or a clearing yield on the SPV's $18B junior tranche, whichever comes first

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
OpenAI most-capable tool-use status	**Paused for training, evaluation and tool-use inference on OpenAI's most capable models (the September pause after an internal agent reached the public internet through DNS, W39); GPT-6.1 Astra withheld (48 days). GPT-6 Cyber shipped to Daybreak Red organisations only, rated Critical**	Paused for training, evaluation, and tool-use inference on OpenAI's most capable models (the September pause after an internal agent reached the public internet through DNS, W39); GPT-6.1 Astra withheld (41 days)	→	A first-party statement that the pause has been lifted

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
TrendForce 2027 blended HBM ASP (average selling price) forecast, year on year	**+121% (not revised); the week's datapoints were supply-side and demand-side at once: Micron reported to have halted construction of its $3-4.5B Taoyuan Fab 11E (HBM / DRAM, originally 2028) over a labour dispute (Bloomberg via Taipei Times, grade 3); CXMT HBM3E yields reported at 20-30% (grade 3); Qualcomm's handset chips -26% on input-cost inflation, memory among the costs named; Sam Altman's PC-DRAM remark and retraction; no per-gigabit contract price disclosed**	+121% (not revised); the suppliers printed around it: SK hynix 80% operating margin with long-term agreements covering more than half of 2027-28 supply for 14 customers and HBM4 overtaking HBM3E in Q4; Samsung HBM4 more than triples to 60% of HBM with 2027 sold out; CXMT HBM3E at 20-30% yields (Reuters, grade 3); no per-gigabit contract price disclosed	→	A revision below +80% or a contracted per-Gb price disclosed by a supplier

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
PJM Reliability Backstop Procurement (RBP) effective date	**Refiled October 29 in ER27-330; requested effective January 15, 2027 with an order requested by January 29, 2027; bid window March 1-28, 2027; comments due November 30. No FERC action this week. Federal context changed: Senate cloture on S. 5653 (which would bar shifting incremental grid costs of ≥20 MW loads to other customers) is set for 5:30 pm November 9 at 20-25% odds (Rapidan via Axios), and Sen. Heinrich, expected Energy chair in January, published a mandatory-cost-allocation fact sheet November 5**	Refiled October 29 in a new docket (ER27-330); requested effective January 15, 2027 with an order requested by January 29, 2027; bid window March 1-28, 2027; comments due November 30; LSE collateral at 40% of discounted payment obligation (~$686k per MW); costs allocated to large-load growth through 2028/29; target capped at the UCAP equivalent of distributor-reported large-load growth	→	A PJM refile accepted without suspension, or the suspension extended past February 2027

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
PJM data-center load cancelled or delayed since the capacity-auction inputs were set	**4,017 MW, per the Independent Market Monitor (no new figure); the Monitor's 3,055 MW adjusted shortfall against PJM's 6,831 MW in ER27-330 stands. Outside PJM, Jefferies estimates ~13 GW of announced Texas projects at risk under the governor-elect's pledged pause (via King & Spalding and Bisnow; grade 4), and ERCOT's audit reports are due December 10**	4,017 MW, per the Independent Market Monitor (no new cancelled-load figure); the Monitor's adjusted backstop shortfall of 3,055 MW against PJM's 6,831 MW in the ER27-330 affidavit implies roughly 3,800 MW of PJM's forecast that the Monitor treats as cancelled or delayed (house arithmetic; within the 4,017 MW figure)	→	The Market Monitor's adjusted shortfall falling to zero, or a reversal that puts cancelled load back into PJM's forecast

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Nvidia InferenceX Vera Rubin submission	**No new InferenceX result; the W43 'official preview' of Vera Rubin NVL72 (1.81-2.03x tokens per TCO dollar at 100 tokens/s per user, modeled TCO, no valid power telemetry) stands. Rubin volume is now contracted: SpaceX's lease covers ~$47.3B of Rubin equipment delivering in H1 2027 (Bloomberg, grade 3)**	No new InferenceX result; the W43 'official preview' of Vera Rubin NVL72 (1.81-2.03x tokens per TCO dollar at 100 tokens/s per user, modeled TCO, no valid power telemetry) stands. Rubin hardware moved: Microsoft docked its first Vera Rubin rack, Oracle made Vera Rubin NVL72 GA on OCI with multi-gigawatt plans	→	A Rubin result on InferenceX for a modern MoE model with cost per million tokens

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Largest model a new handset platform claims to run locally	**30B-parameter mixture-of-experts on Snapdragon 8 Elite Extreme Gen 6, vendor claim (no independent result). Qualcomm's handset chip revenue fell 26% on memory costs, which is the device-side bill for the memory the claim depends on**	30B-parameter mixture-of-experts on Snapdragon 8 Elite Extreme Gen 6, vendor claim (no independent result)	→	An independent on-device token-per-second result for that model class

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Lumentum AI optical components: sold-out horizon	**Early 2029 (unchanged); the object is now under a definitive agreement: Marvell will acquire Lumentum for 4.500 shares per share ($115B implied EV, 23.9% premium; close H1 2028). Lumentum FQ1 revenue $1.26B (+136.5%) with a second NPO / CPO win backed by >$5B of commitments; Coherent reports no cancellations with coverage into 2028 (p143, the W43 bet on exactly that, resolved as a hit); AOI cut its end-2026 capacity target on power and construction, a supply-side cut**	Early 2029 (unchanged); shares +8% on October 26 (Coherent +6%) on a StoneX note, reversing part of the October 22 selloff; Bloomberg reports Marvell nearing an ~$85B deal for Lumentum (single source, grade 3, scored noise); Corning's generative-AI optical revenue doubled and is 'significantly outpacing GPU unit growth'	→	The horizon moving in by two quarters, or a second optics vendor reporting a cancellation

THRESHOLD VALUES ARE THE POINTS WHERE THE READ FLIPS. CROSSINGS ARE FLAGGED IN THE ISSUE BODY WHEN THEY HAPPEN.

PREDICTIONS

What we expect next.

5 falsifiable, time-bounded predictions. Each carries a confidence (1–99%, never 50%), a deadline, and a specific signal we'll watch. Future issues score them hit / miss / partial.

PREDICTION 01 CAPITAL

62%

BY JANUARY 31, 2027

By January 31, 2027 a second hyperscaler or chip vendor (Amazon, Microsoft, Google on a new site, NVIDIA, AMD or Broadcom) is disclosed in a landlord's or lessor's SEC filing as providing credit support, a guarantee or a lease-assumption right for a frontier lab's data-centre or compute lease with an unaffiliated landlord, that is, the Google-TeraWulf agreement (hyperscaler side) and NVIDIA's August 2026 residual-value guaranties on OpenAI's SB Energy leases (vendor side) are the first two instances of a structure in which the seller of compute lends its credit so a third party will finance the buyer, and a further instance from a third counterparty follows within the quarter.

TRIGGER: Hit if, by the deadline, a Form 8-K, 10-Q, 10-K or prospectus on EDGAR from a landlord, lessor or developer (any company other than the lab and the guarantor) discloses an executed agreement under which a hyperscaler or accelerator vendor guarantees, backstops, agrees to assume, or provides credit support for a frontier lab's (OpenAI, Anthropic, SpaceX SI / xAI, Meta Superintelligence as tenant of a third party, Mistral, or any lab with more than $1B of annual revenue) lease or capacity contract of $1B or more. A hyperscaler guaranteeing its own affiliate's lease, a vendor equity investment in the landlord, a vendor supply commitment, or a report of negotiations without a filed agreement is a miss on the thesis leg; a disclosed agreement that names 'a hyperscaler' or 'a strategic partner' without identifying it is a partial. Google's existing Fluidstack backstops (2025), NVIDIA's August 17, 2026 residual-value guaranties on OpenAI's SB Energy leases (8-K accession 0001045810-26-000069, outside the graded packet) and the TeraWulf agreement itself do not count; a new agreement by Google or NVIDIA with a different landlord or lab does. The bet sits above even because the mechanism has now appeared on both sides of the trade (Fluidstack twice, outside the graded packet, and TeraWulf once on Google's book; SB Energy on NVIDIA's), because Anthropic's S-1/A shows ~$86B of neocloud offtake to Nscale and Lambda whose landlords and lenders face the same tenant-credit question, because IFR reports five pulled data-centre IPOs and developers offering tenants equity, which this issue reads (its inference, not IFR's) as demand for exactly this enhancement, and because Amazon and Microsoft each have a lab counterparty and a silicon roadmap to attach to the guarantee as Google attached TPUs. It is held below 70 because a guarantee is a contingent liability a hyperscaler's board and auditors must approve, because the SpaceX lease (the next largest structure) is being sold without one, and because twelve weeks is a short window for a landlord filing.

PREDICTION 02 HARDWARE

40%

BY JANUARY 31, 2027

By January 31, 2027 Nvidia discloses in a first-party document (10-Q, 10-K, 8-K or press release) a direct investment, prepayment, capacity-reservation advance or equipment financing of $1 billion or more to an outsourced assembly-and-test provider (ASE / SPIL, Amkor, PTI or another OSAT) for advanced packaging, that is, the Reuters report that Nvidia is providing 'billions' to expand CoWoS capacity at SPIL becomes a disclosed line in Nvidia's own filings within one reporting cycle.

TRIGGER: Hit if Nvidia's Form 10-Q for the quarter ending late October 2026 (due by early December), any 8-K, or a press release on nvidianews.nvidia.com by the deadline discloses a payment, prepayment, investment or financing commitment of $1B or more identified as going to a named OSAT or to 'packaging' or 'assembly and test' capacity at a third party; an aggregate 'prepaid supply' or 'supply agreements' figure that is not attributed to packaging or to an OSAT is a partial if it rises by $3B or more quarter on quarter (the W43 p141 condition) and a miss otherwise; a disclosure by the OSAT alone (ASE's or Amkor's filings naming Nvidia) is a partial; silence is a miss. The bet sits at 40 because Nvidia's W43 10-Q already itemised $1.1B of capacity-reservation advances into test and $1.3B into Amkor's interposers, so the disclosure practice exists and the Reuters report is specific about the counterparty, and because KGI's CoWoS projection (120k to 185-190k wafers a month) is the scale at which Nvidia's own supply depends on the OSAT layer. It is held below even because the Reuters item is grade 3 with no first-party confirmation, because Nvidia's prepayments are typically disclosed in aggregate rather than by counterparty, and because 'billions' may be structured as purchase commitments that never appear as a line. A hit confirms the hardware lens's reading that packaging, not power, binds the 2027 build; a miss leaves it a reported intention.

PREDICTION 03 NETWORKING

66%

BY MARCH 1, 2027

Equinix's fourth-quarter 2026 results, reported by March 1, 2027, show interconnection revenue growing faster year on year than colocation revenue for a second consecutive quarter (Q3 2026 was +10% against +7%), that is, the Metcalfe test the framework's networking hypothesis sets (interconnect growth outpacing capacity growth across two consecutive quarters) is met on the operator's own segment lines.

TRIGGER: Hit if Equinix's Q4 2026 earnings release or supplemental on its investor-relations site (or the 8-K exhibit on EDGAR) by the deadline reports a year-on-year growth rate for interconnection (or 'interconnection and other' if the line is combined) that exceeds the year-on-year growth rate for colocation (or 'colocation and related') on a reported or constant-currency basis, with both figures stated by the company; a gap of less than one percentage point is a partial; interconnection growth at or below colocation growth is a miss, as is a segment restatement that removes the comparable lines, in which case the resolver uses the nearest disclosed interconnection metric (cross-connect count growth against cabinet-equivalent growth) and notes the substitution. The bet sits at 66 because Q3's gap was three points, because Fabric One's public preview (free to 3 Gbps; AWS Interconnect live in two metros; GA 2027) is an interconnection product shipping into the quarter, because Lumen's hyperscaler fibre deals above $13.5B and AI backbone traffic up 1,000% since April say the inter-site traffic that cross-connects monetise is growing faster than the cabinets, and because colocation growth is capped by delivered capacity in a power-constrained year. It is held below 75 because one quarter's three-point gap is inside the range of ordinary quarterly variation at Equinix, because colocation pricing has been rising on scarcity and could lift that line, and because currency and a hyperscale-lease timing effect can move either line by a point. A hit converts the networking hypothesis's two-quarter test from strained to met at the operator layer; a miss sends it back to the component layer, where this week's prints already passed it.

PREDICTION 04 SOFTWARE

48%

BY DECEMBER 31, 2026

By December 31, 2026 Anthropic states in a first-party document (an RSP update or capability report on anthropic.com, an S-1 amendment or final prospectus on EDGAR, or a system card) that it has met, or determined that its models have crossed, the AI R&D acceleration threshold in its Responsible Scaling Policy, that is, the S-1/A's 'may imminently' and METR's measured ~1.9x become an acknowledged threshold crossing before the year ends.

TRIGGER: Hit if, by the deadline, a document published by Anthropic (RSP page or update, a capability or risk report, a model system card, or an S-1 amendment or 424(b) prospectus on EDGAR) states that the AI R&D acceleration threshold, capability threshold or 'ASL' trigger tied to AI-driven R&D acceleration has been met, crossed or triggered, or that the company has moved to the safeguards that threshold requires; a statement that the threshold 'may' or 'is likely to' be met, a statement of new safeguards without reference to the threshold, or a third-party assessment alone (including a further METR report) is a miss on the thesis leg; a statement that the threshold has been met for internal models but not deployed ones is a hit. The bet sits just below even: on one side, the S-1/A risk factor says the company 'may imminently' meet the threshold, METR measured a 1.9x improvement rate (95% CI 1.5-2.8x) with ~80% AI-attributed over eleven weeks of high-access assessment, and an IPO prospectus that must be updated for material developments is a forcing document; on the other, a threshold crossing triggers safeguard obligations a company would rather announce after the IPO prices than before, the RSP's threshold definition leaves the company judgement about measurement, and 'may imminently' is language a filer can carry for months. Either outcome informs the first hypothesis: a hit is the strongest acceleration datapoint the framework will have had, from the lab's own governance document; a miss says the lab's measurement lags the external one by at least a quarter.

PREDICTION 05 POWER

57%

BY FEBRUARY 18, 2027

Within thirty days of taking office (by February 18, 2027), Texas Governor Gina Hinojosa issues an executive order, proclamation or directive that pauses or suspends state approvals (PUCT / ERCOT interconnection or energization approvals, or TCEQ permits) for new data centres above a stated megawatt threshold pending legislation, that is, the pledged moratorium becomes a signed instrument rather than a campaign position, extending the outgoing governor's audit pause rather than letting it lapse.

TRIGGER: Hit if, by the deadline, an executive order, proclamation or directive letter signed by the governor and published on gov.texas.gov (or filed with the Texas Secretary of State) directs the PUCT, ERCOT or TCEQ to pause, suspend or withhold approvals, interconnection or energization decisions, or permits for new data-centre loads above any stated threshold (25 MW, 50 MW, 75 MW or otherwise) for a defined period or pending legislative action; a directive that only continues the existing audit without pausing approvals, a request to the legislature without an instrument, or a pause limited to state tax exemptions is a partial; no instrument by the deadline is a miss. The bet sits above even because the pledge was specific (an executive moratorium on all approvals until the 90th Legislature passes the 'Texans First' package), because the outgoing governor's August PUCT / ERCOT audit pause and September TCEQ permit halt have established that these agencies will act on a governor's letter, because the ERCOT audit reports due December 10 give the incoming administration a document to act on, and because the Texas House flipped to the governor's party, lowering the political cost. It is held below 65 because the Texas Senate remains 19-12 Republican and PUCT appointees need two-thirds Senate confirmation, because Jefferies' ~13 GW at risk is the kind of number that draws a negotiated narrowing before inauguration, because the outgoing governor controls the agencies until January 19 and may lift or extend the pause himself (which would change what 'extending' means), and because governors-elect frequently convert moratoria into audits once in office. A hit puts the largest US data-centre market under an executive pause and gives the fifth hypothesis its permitting leg; a miss says the ballot-box signal of November 3 did not survive contact with the state's interconnection queue.

TRACK RECORD

Every prediction ever made, scored against its own written trigger when the deadline passes; ambiguity resolves against us.

Predictions made	**156**	Resolved (hit / partial / miss)	**85 (37 / 20 / 28)**
Hit rate (partial = half)	**55%**	Brier score (0 = perfect)	**0.192**

Bold (<55%): 7 resolved, 57% hit rate vs 40% mean confidence · Core (55-80%): 74 resolved, 53% hit rate vs 65% mean confidence · High-conviction (>80%): 4 resolved, 88% hit rate vs 83% mean confidence

WATCHLIST

On the radar.

7 catalysts in the next 7–14 days that would change the read materially. Watching these tells us whether the thesis is strengthening or weakening.

NOVEMBER 9, 5:30 PM ET

WATCH 01

Senate cloture on H.R. 2347 / S. 5653 (permitting vehicle with the ≥20 MW covered-load cost-allocation bar)

Needs 60 votes at 20-25% odds (Rapidan via Axios); failure sends the same cost-allocation principle to a January Senate with Heinrich at Energy, which is the federal leg of PJM's RBP design and the context for p149.

NOVEMBER 10

WATCH 02

CoreWeave Q3 2026 results

Guided interest expense of $860-940M against adjusted operating income of $200-260M, 2032 notes above 16% and CDS ~1,155 bp; the print is the within-issuer test of the spread lever and the first look at how the ladder's bottom rung funds a $104B backlog Bernstein says needs ~$85B more.

WEEK OF NOVEMBER 9 THROUGH DECEMBER 5

WATCH 03

Anthropic roadshow and pricing (p145)

The S-1/A's price range is still blank; CNBC reports investors see $1.0-1.25T as more likely than $1.5T, which is the partial band for p145; the prospectus must also update any RSP acceleration-threshold determination (p155).

NOVEMBER 20 AND NOVEMBER 30

WATCH 04

Comment deadlines: NERC RR27-1 large-load registry (November 20) and PJM ER27-330 RBP refile (November 30)

The two dockets that decide whether registration and backstop procurement for large loads become federal policy on PJM's timetable; the Market Monitor's next adjusted-shortfall figure will appear in the ER27-330 record and moves the cancelled-load lever.

DECEMBER 1 (ARGON); DECEMBER 4 (SWE-2 FLASH)

WATCH 05

Two dated price steps: Gemini 4 Argon reverts from the introductory $2 / $10 to $4 / $20; SWE-2 Flash moves from free in Devin to $0.10 / $0.50 per million tokens

Argon's step would be the first list increase inside the $2 / $10 band in the series and takes the band to three labs unless Google extends; Carbon's API status by December 15 decides p147 (38%). Separately, SWE-2 Flash's first paid price tests whether Cognition publishes the withheld weights alongside it (the Model Pulse tracks the gate).

DECEMBER 10

WATCH 06

ERCOT data-centre audit reports to the PUCT

The first official count of how much announced Texas load the audit treats as real, the document the incoming governor's pledged pause would act on (p156), and the Texas analogue of the PJM Monitor's cancelled-load figure.

NEXT SOURCED PRINT

WATCH 07

Broadcom five-year CDS absolute level

CNBC reports a 49% rise over the month without a level; applied to the W42 base of 133-136 bp that is ~200 bp, the lever's threshold, but the W42 print and Deutsche Bank's month-ago base may differ, so the implied level is house arithmetic. A sourced quote either confirms the market is pricing Broadcom as a residual-value financier or returns the lever to flat.

METHODOLOGY AND AUTHORSHIP

How this brief is built.

Compiled from public announcements, SEC filings, earnings transcripts, and official lab and vendor publications. Every quantitative claim is graded on source quality. Every prediction is falsifiable, time-bounded, and scored hit / miss / partial / pending in future issues.

SIGNAL SCORE RUBRIC

5 / 5	SEC filing or audited disclosure. Multi-source independent confirmation. Operational, not aspirational.
4 / 5	Earnings-call disclosure or primary lab/vendor announcement. Two or more independent sources.
3 / 5	Single primary source. Reasonably consistent with sector data.
2 / 5	Analyst report or secondary press only. Or single primary source with credibility caveats.
1 / 5	Rumor, social media, single non-primary source. Or contradicted by alternate primary sources.

ANYTHING GRADED 2 OR BELOW IS FLAGGED AS NOISE.

PREDICTION OUTCOMES

HIT	The specific observable outcome occurred by the deadline.
MISS	The deadline passed and the outcome did not occur.
PARTIAL	A meaningfully similar outcome occurred but not the literal wording.
PENDING	Deadline not yet reached.

HIT RATE OF 65-75% IS THE TARGET. ABOVE 80% MEANS TOO CONSERVATIVE; BELOW 50% MEANS THE FRAMEWORK IS WRONG.

AUTHORSHIP

Brian Letort

BrianLetort.AI · Independent analysis. All sources public. Not investment guidance.

THIS ISSUE

Issue 29

Week 45 of 2026 · November 7, 2026

WEB

brianletort.ai/industry

Past issues, the working framework, and the LLM Evolutionary Tree companion.