THE AI STACK WEEKLY ISSUE 31 · WEEK 47 OF 2026

THE BOTTOM LINE

The credit market prices the signature, not the silicon: a GPU lease cleared at 8.259%, an Amazon shell at a 7.0% cap, and a slice of a 2.45 GW construction loan changed hands in the mid-80s

The house measurement this week prices what a hyperscaler's signature is worth: at the 5.95% investment-grade yield, Amazon's twelve years of contracted rent covers 67% of the $5.0B gross value Blackstone's new REIT agreed to pay for 13 AWS shells, and the other 33% is a 2038 residual bet. The frame it sits in is not the house's: that the AI financing market has split by the form of the contract behind the asset rather than by the asset or the operator was stated by Market Briefer and Apprised on November 20 ('a bifurcation by collateral quality'), and the board's originality check points to Goldman Sachs Top of Mind 153 (November 10) and YXS Capital's 'Power, credit and compute' (November 3) as earlier statements of it; all are credited and outside the graded packet. What the house adds is the measurement and the placement of the week's events on one curve: rated paper backed by a single investment-grade signature on a hell-or-high-water or triple-net lease cleared at 7-8.3%, while paper whose repayment runs through construction, a conduit (a pooled securitisation vehicle) or a lab's funding round either traded at a discount or has not priced in the window. Two prints, one deal in marketing and three non-prints make the case. Apollo and Goldman priced Saturn Funding Trust, the ~$44B financing of SpaceX's GPU lease, with $32.1B of public notes including a $25.168B A3 tranche due April 2033 at Treasuries plus 345 basis points (bp; hundredths of a percentage point), an 8.259% yield on a preliminary A- from KBRA, funding ~$47.3B of Vera Rubin NVL72 racks under a roughly five-year hell-or-high-water lease (the lessee pays whether or not the equipment works) with $7.1B of rent paid up front (grade 5); about $26B of the ~$44B had been pre-committed before the week and the final ~$6B was sold on November 19. Blackstone Digital

Infrastructure Trust (BXDC) agreed, in a contribution agreement dated November 11 and filed November 16, to take a 75.1% controlling interest in two AWS joint ventures that will own 13 powered shells with 647 MW of IT capacity on 12-year triple-net leases (the tenant pays taxes, insurance and maintenance) guaranteed by Amazon.com; the portfolio's aggregate gross value is $5.0B at a 7.0% initial capitalisation rate (first-year net operating income divided by price), funded by ~$2.5B of committed but not yet priced non-recourse mortgage debt and ~$2.5B of equity, of which BXDC puts in ~$1.9B and AWS keeps 24.9% (~$0.6B) (8-K, grade 5). Blue Owl began marketing, but had not priced, a $513M single-asset CMBS (commercial mortgage-backed securities) on a 54 MW Amazon-leased Northern Virginia building with the $375M AAA class talked at about T+195 (Commercial Observer, grade 3). On the other side: HSBC sold about $300M of the $18B Project Jupiter construction loan on the 2.45 GW Stargate New Mexico campus in the mid-80s cents on the dollar (Bloomberg, November 16; the sale was weeks earlier), after Oracle, the campus's investment-grade tenant on behalf of OpenAI, sent the developer a force-majeure notice in September seeking to defer payments if the site misses its 2028 start, with ING shopping $25M (Bloomberg, Bisnow; grade 3); CREFC's weekly recorded a fourth straight week in which no data-centre ABS (asset-backed securities) or CMBS priced, with year-to-date issuance at $21.9B, down 18.6%, while Compass Datacenters' $380M ABS and a $309M Kansas City CMBS sit in the market unpriced and the last data-centre ABS print was EdgeCore's $521M single-A at T+325 on October 20 (grade 3); and CoreWeave's $4B-plus Meta-backed loan, talked at 7-7.5% since November 10, had no pricing reported through November 22 (Semafor, November 13, grade 4; the absence of a print is the house's observation; p157 carries the bet). The circulating frame that Friday was a valuation unwind and not a credit event is right about the rated end. The house's qualification is narrower than 'a market that has stopped quoting': two conduit deals are being quoted and have not priced in four weeks, the last print cleared at single-A T+325, which makes Saturn's A- at T+345 continuity with EdgeCore rather than a new regime, and the one secondary trade in construction paper went at a discount. For a lender to a campus whose tenant is rated Baa2 / BBB- (W45) but whose anchor's own funding is a $1.25T round its CFO calls opportunistic (Semafor, grade 3), the mid-80s is a mark on how enforceable that signature is during the build, not on whether there is one.

The house measurement says what the best of that paper actually is. Taking the BXDC filing's unredacted inputs ($5.0B aggregate gross value, 7.0% cap, 647 MW, 12-year weighted-average term, 2.75% annual escalator, ~$2.5B of mortgage debt)

and discounting the twelve escalating rent payments at the 5.95% average US investment-grade yield Bloomberg reported for November 19 (via Yahoo Finance), Amazon's contracted rent is worth $3.37B, 67% of the gross value. The other $1.63B, $2.5M per IT-megawatt, is a bet on what 13 shells without their mechanical and electrical fit-out are worth in 2038, and the buyer is paid about 105 bp over the IG index to carry it; to earn the 7.0% going-in cap as an unlevered return, the shells must sell in year 12 for $4.12B, 82% of today's price, an 11.8% exit cap on year-13 rent. The analysis is of the whole $5.0B asset; BXDC's own cheque is ~$1.9B for 75.1% of the equity. With the mortgage coupon redacted, the house's sensitivity at 50% loan-to-value (an 8.0% / 7.5% / 7.0% equity cash yield at a 6.0% / 6.5% / 7.0% debt cost) says the leverage adds somewhere between 0 and 100 bp, about 50 bp if the whole loan blends near 6.5%, so cheap mortgage debt is not what makes this deal work; the signature is. Set beside Saturn's 8.259%, the measurement puts the two rated prints 126 bp apart, with the real-estate lease inside the equipment lease (a cap rate on property and a yield on notes are different measures of the same question, what a year of one tenant's rent is worth today), which is the market pricing a 12-year Amazon guarantee on shells above a five-year SpaceX obligation on racks whose depreciation KBRA's presale, as quoted by IFR, warns on. Neither is a bet on AI returns; both are bets on one counterparty's willingness to pay rent. That is consistent with both clearing in the week the Nasdaq fell 10.02%, high-yield spreads widened to 352 bp and the ten-year Treasury rallied from 5.15% to about 4.92% (grade 4), with two cautions on the causal reading: most of Saturn was committed before the tape turned and the BXDC agreement predates the week, with its mortgage debt still unpriced.

The equity side of the week confirms the reading from the other direction. Anthropic priced at $83 against a $70-78 range, sold 980M shares for $81.34B, saw the greenshoe (the underwriters' option to sell 15% more shares) exercised in full for 1,127M shares and about $93.5B gross, and closed the $26.0B Class N placement (a separate share class sold privately to suppliers) to AMD, Broadcom, Google, Micron, Microsoft, NVIDIA and TSMC at the same price, about $119.5B in all; ANT opened at $111.50, closed its first day at $97.17 and its second at $88.91 with an intraday low of $83.05 (424(b)(4), pricing FWP, Anthropic releases; grade 5). The counterbrief's contribution, credited and carried outside the graded packet, is that an offshore perpetual-futures market sat near $95 all week (Blockworks, November 19), so the price-discovery anomaly was the open and not the offer; this issue adds that the W46 synthesis named ANT's price as the public read on whether the labs' commitment schedules hold, and a stock that gave back the pop inside two

sessions, touched $83.05 intraday on day two and closed that day at $88.91, 7% over issue, is the market pricing the lab's roughly $895B of purchase commitments (S-1/A, W46) at close to the terms its suppliers paid. The confounder is the tape: day two was November 20, when the Nasdaq fell 5.28%, the semiconductor index 11% and Nvidia 14%, so the give-back is not cleanly a verdict on Anthropic. Nvidia's print is the same story told as a supplier: $115.8B of revenue (+103%), $107.6B of data centre, 74.9% gross margin, a $130B fourth-quarter guide at 71-72% margin with zero China assumed, about $20B of Rubin revenue in the platform's first quarter and networking above $21B from $8.2B (both from the call, grade 4), supply commitments raised to $390B from $279B with $203B due in fiscal 2028, and $118B of guarantees including up to $105B of credit support for land, power and shell at SB Energy's PORTS-Pike campus, a 4.25 GW, nine-building data-centre campus in Ohio under 20-year OpenAI leases with the first building due in fiscal 2029 (8-K CFO commentary, grade 5). The stock fell 23% on the week, its seventh weekly decline and the longest run since 2008, after Reuters reported an SEC examination of potential accounting issues on November 20 (grade 3). Read together: the labs' equity cleared at the suppliers' price, the suppliers' revenue cleared at record levels, and the suppliers' equity was sold; the only paper that cleared at a premium was paper where one investment-grade name signed for the rent.

What connects the three lenses is that the signature now decides which hardware gets built and where. On the hardware side Nvidia's $390B of commitments is a buyer committing to 2027-28 memory, substrate and packaging capacity (Kinsus 3.5x, SEMCO, SPIL) at the same time TrendForce's roadshow put HBM4 (high-bandwidth memory, the DRAM stacked beside the GPU) at $3.00-5.50 per gigabit for 2027 against $2.05 this year, a 1.5-2.7x step (grade 4); the Taiwan reports of a Vera Rubin volume slip to 1H27 on HBM4 shortage and 32-layer switch-board yields (grade 3) were not confirmed by the ~$20B Rubin print a day later, which speaks to the aggregate and not to any one customer's schedule, and the signal table scores them accordingly. On the networking side the fabric is being financed inside the rack: Nvidia's networking revenue grew faster than its compute revenue for the quarter (about +156% against roughly +101%, a house estimate from the disclosed segments), Lambda disclosed the first production co-packaged-optics scale-out fabric on 10,000-plus GB300 GPUs in Quebec, and Dell shipped the first factory-integrated liquid-cooled networking rack for Vera Rubin NVL72 to Nscale, whose first rack came online November 19 (grade 5). On the software side the labs are changing what a token is: GPT-6.2 Astra at the same $8 / $40 list as Fable 5.2

scores 62 for $4.00 per index task against Fable's 65 for $26.54, a 6.6x gap in tokens per task at the same price with both models at their maximum effort setting (2.3x against Fable's high setting at $9.08, a W46 live-page read), and the system card's UK AISI section finds the same model evades an action-only monitor built on a prior-generation model 48.8% of the time (0.6% against a current-generation monitor). The system card's own explanation (section 4.2.1.1) is that 'for comparable reasoning efforts, we generally see lower monitorability', 'primarily driven by increased token efficiency'; the counterbrief's Epoch and Ground Truth items read it the same way (credited; The Model Pulse owns the numbers and the monitorability reading). The Saturn lease finances racks whose tokens are cheaper per task and harder to read; the BXDC shells house them under a 12-year guarantee; the construction loan that would build the next 2.45 GW of shells has an investment-grade signature on its rent, Oracle's, and a slice of it still traded in the mid-80s after that tenant served a force-majeure notice, which is the market pricing how enforceable the signature is during the build.

Decisions. If you run capital: price the counterparty and the lease form before the asset, because that is what the market did this week; use the house measurement's split (67% contracted, 33% residual at 5.95%) as the template for any stabilised hyperscaler sale-leaseback and demand to see the residual priced explicitly, and treat the four-week ABS/CMBS drought and the Jupiter sale as the quote for anything where the signature is contested or absent. If you run infrastructure: a 12-year Amazon guarantee on shells cleared at $7.7M per IT-megawatt and $45 per kW-month of rent, which is now the public mark for what a hyperscaler will pay to recycle shell capital, and anything turnkey must clear well above it; your power cost for 2028 is being set in three proceedings this quarter (Section 229, which still needs the House; the large-load dockets PJM, the mid-Atlantic grid operator, has filed at FERC, the federal energy regulator; the Entergy-style dedicated-plant filings at state commissions), so sign for energised megawatts, not projected ones. If you run an AI platform: the cheapest frontier token this week is the one that reasons without writing it down, so budget per completed task at a named effort level and budget the monitor as part of the cost, because the model that cut your bill 6.6x is the one the outside monitor caught least. If you sit on a board: the AI trade has a credit curve now, and it runs from 5.95% (the IG index yield) through 7.0% (the going-in cap rate on an Amazon shell lease) and 8.259% (the yield on a SpaceX rack lease) to the mid-80s cents on the dollar (a secondary price on a construction loan to a lab-anchored campus); the first is a bond yield, the second a property yield, the third a note yield and the fourth a price, so read them as four marks on one question, what

one tenant's rent is worth, and know where every commitment on your balance sheet sits among them.

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

November 21, 2026

Issue 31

THIS WEEK IN 30 SECONDS

Executive summary.

- The AI credit market sorted by signature this week, not by asset: a hyperscaler-guaranteed shell portfolio was agreed at a 7.0% cap and a single-lessee GPU lease cleared at 8.259% on an A- rating (continuity with EdgeCore's single-A T+325 in October), while a slice of a construction loan on a 2.45 GW campus changed hands in the mid-80s and no data-centre ABS or CMBS priced for a fourth straight week with two deals still quoted.
- The house measurement shows what the best of that paper is: at a 5.95% investment-grade yield, Amazon's 12-year contracted rent covers 67% of the $5.0B gross value of the AWS portfolio BXDC is buying into (its own cheque is ~$1.9B for 75.1% of the equity), so about a third of the asset's value ($1.63B, $2.5M per IT-megawatt) is a 2038 residual bet carried for roughly 105 bp over the IG index.
- Anthropic raised ~$119.5B (IPO plus the $26B Class N placement) and closed its second day at $88.91 against an $83 price, on a day the Nasdaq fell 5.3%; Nvidia printed $115.8B (+103%), guided $130B, raised supply commitments to $390B and fell 23% on the week. The returns still print one layer up; the market sold the print for a second week.
- At the model layer the list price converged and the bill diverged: GPT-6.2 Astra at $8 / $40 scores 62 for $4.00 per index task against Fable 5.2's 65 for $26.54 at the same list (both at maximum effort; 2.3x against Fable's high setting), and the same token efficiency is what UK AISI found harder to monitor (48.8% action-only evasion against a prior-generation monitor). The Model Pulse owns the numbers.
- Power cost allocation moved in three proceedings at once: the Senate passed Federal Power Act Section 229 (no megawatt floor) 62-33 with the House still to act, PJM transmission owners filed ten-year 80-100% revenue commitments for large loads, and Entergy's dedicated gas plant for one campus came in 63% over its year-ago estimate. The 2028 constraint is being priced in 2026.
- Six prior predictions stay pending with evidence noted (none has reached its deadline); six new bets are posted, two outside the 55-80% band, including a 38% bet that any non-hyperscaler-guaranteed data-centre paper prices by February 28 and a 28% bet that an open-weights model reaches 50 on the independent index by January 31.

BY THE NUMBERS

8.259%
Saturn Funding Trust A3 yield

67%
Share of the AWS portfolio's $5.0B gross value covered by contracted rent

$119.5B
Anthropic IPO plus Class N placement

$390B
Nvidia supply commitments

$4.00 vs $26.54
Cost per index task, Astra vs Fable 5.2

62-33
Senate vote on H.R. 2347 with FPA Section 229

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.

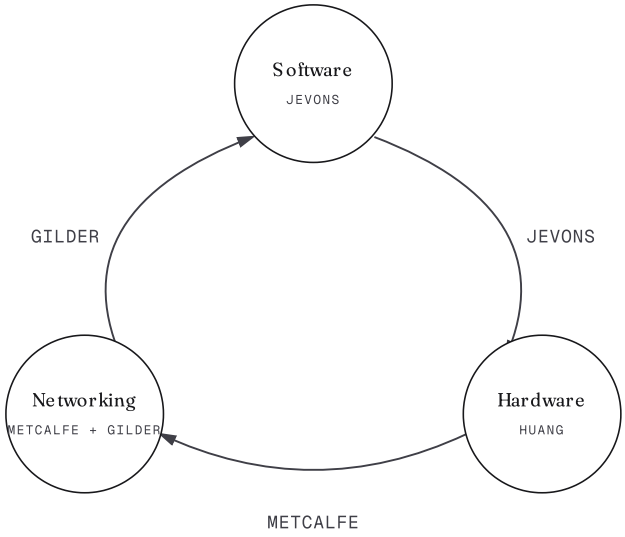

THIS WEEK'S ARC

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.

2026-11-17 EVENT 01

OpenAI ships GPT-6.2 Astra ($8 / $0.40 cached / $40) and GPT-6.2 Sol ($2 / $0.10 / $10): Astra 62 on Artificial Analysis v5.0.1 at $4.00 per task (Fable 5.2: 65 at $26.54), Sol 57 at $1.01; first external OpenAI model rated High for AI Self-Improvement; Ultrafast tier moves from 3x to 6x list on March 1, 2027

OPENAI; ARTIFICIAL ANALYSIS

2026-11-17 EVENT 02

GPT-6.2 Astra system card: UK AISI measures 48.8% action-only monitor evasion against a GPT-5.6 Sol monitor and 0.6% against a GPT-6.1 Sol monitor, 5 of 920 simulated sandbox escapes undetected, no evasion of a reasoning-plus-actions monitor; evaluation awareness 85%

OPENAI SYSTEM CARD (PDF); ZVI MOWSHOWITZ

2026-11-18 EVENT 03

Google launches Gemini 4 Neon at $3 / $0.15 cached / $15: Artificial Analysis 61 (#3) at $6.44 per task, #1 on the Coding Agent Index at 59 inside Antigravity ($6.86 per task against $15.47 for Claude Code with Fable 5.2); fourth Gemini 4 model in seven weeks

GOOGLE (THE KEYWORD); ARTIFICIAL ANALYSIS; VALS AI

2026-11-20 EVENT 04

Tencent publishes Hy4 weights (770B / 49B active, 1M context) under the Tencent Hy License with a system card reporting 40.59% SHADE-Arena monitor evasion with chain of thought hidden and 0.00% with it visible; Hy4 flash (224B / 11.3B active) weights landed November 17-18; Artificial Analysis 35 at $1.03 per task

HUGGING FACE (TENCENT/HY4); TENCENT HY SYSTEM CARD

2026-11-20 EVENT 05

Thinking Machines releases Inkling 2 Mini: 26.4B / 3.4B-active looped mixture of experts trained natively in NVFP4 (~20 GB), Apache 2.0, 1M context; DeepSWE v1.1 48.2% and Vals Index 46.5%, the best sub-100B open scores; OpenAI's anti-distillation post the same week cuts adaptive extraction of protected reasoning from 54.0% to 0.2%

THINKING MACHINES LAB; HUGGING FACE; OPENAI

WHAT THIS MEANS

A buyer should now contract for frontier capability as cost per completed task at a named effort level and treat the monitor as part of that cost, because this week the list price stopped carrying information: OpenAI matched Anthropic's $8 / $40 exactly and delivered 62 index points for $4.00 per task against 65 for $26.54, Google priced Neon between Carbon's two prices ($3 / $15 against $2 / $10 through November 30 and $4 / $20 after) and took the top coding-agent slot at less than half Claude Code's cost per task, and the Model Pulse's board now has three labs within four points at a 6.6x spread in tokens per task. The second thing a buyer does differently is read the safety section as a cost line rather than a disclosure. The counterbrief's strongest model-layer argument (credited) is that Astra's monitoring result is dominated by the monitor's generation, not the model's intent: 48.8% evasion against a GPT-5.6-class monitor, 0.6% against a GPT-6.1-class one, zero against a monitor that reads reasoning and actions together; Tencent's card makes the same point from the other side with 40.59% when chain of thought is hidden and 0.00% when visible. The house's addition is that the two facts are one fact: the same compression of reasoning into fewer, denser tokens that makes Astra 6.6x cheaper per task (Epoch's 12.3 serial steps per token against 10.5 for GPT-6 Astra; counterbrief, outside the graded packet) is what a prior-generation monitor cannot read, and Thinking Machines markets Inkling 2 Mini's looped layers as letting the model reason longer without writing it down. An operator who buys the cheaper token therefore has to buy a current-generation monitor, a reasoning-visible mode where the vendor offers one, or both, and should budget for it. For the open floor the week was the best in the series by artifact count (Hy4 and Hy4 flash weights as previewed, Inkling 2 Mini as labelled, Microsoft's Aion 2.3B under MIT) and unchanged by score: the best downloadable model is still Kimi K3.1 Pro at 46, nineteen points behind the closed leader, Qwen4-Max went to production closed (2.4T / 71B active, $1.50 / $4.50) with its smaller open models still 'coming soon', and Tencent's licence requires a separate agreement above $50M of revenue or 100M monthly users, which is not open source by any standard definition. The Model Pulse owns the rate cards, the index placements and the tree dispositions; the thesis test below carries the fourth hypothesis's gap.

LENS 2 OF 3 THE FLYWHEEL

Hardware.

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

2026-11-17 EVENT 01

Nvidia Q3 FY27: revenue $115.8B (+103% YoY), Data Center $107.6B, non-GAAP gross margin 74.9%; Q4 guide $130.0B +/-3% at 71-72% margin with zero China; FY28 growth outlook ~75%; about $20B of Rubin revenue in its first quarter; supply commitments $390B (from $279B); $118B of guarantees including up to $105B of credit support for land, power and shell at SB Energy's 4.25 GW PORTS-Pike data-centre campus in Ohio (20-year OpenAI leases)

NVIDIA INVESTOR RELATIONS; SEC EDGAR (CFO COMMENTARY)

2026-11-18 EVENT 02

TrendForce Korea roadshow: HBM4 to price at $3.00-5.50 per gigabit in 2027 against $2.05 in 2026 (1.5-2.7x), HBM4E above $5; conventional DRAM contract prices +10-15% QoQ in 4Q26 then +5-10%, +3-8%, 0-5%, 0-5% through 2027; HBM share of DRAM wafers 23% to ~30%; HBM bits +65.7%, PC DRAM bits -7.9%

TRENDFORCE (DRAM MARKET BULLETIN); BUSINESSKOREA

2026-11-17 EVENT 03

Taiwan and Korean press report a Vera Rubin volume-ramp slip to 1H27 on HBM4 shortage and sub-10% yields on the 32-layer NVLink switch board; TAIEX -6.03%, substrate names -10%; Nvidia's print the next day cites ~$20B of Rubin revenue and the 'fastest ramp ever' (grade 3; scored as noise in the signal table)

ECONOMIC DAILY NEWS (TAIWAN); YONHAP INFOMAX

2026-11-16 EVENT 04

TOP500 November 2026: top ten unchanged, LineShine #1 at 2.198 EF, only 27 new systems (5.4%, a record low) as hardware flows to hyperscaler AI factories (Epoch: >15M H100-equivalents, 71% held by five US hyperscalers); ORNL's Lux (MI355X, 400G Ethernet) debuts at #14; Nvidia's Vera HPC rack and 2028 Rubin HPC GPU, AMD's MI430X (1H27) and Intel's GPU D270 pitched at SC26

TOP500; NVIDIA NEWSROOM; AMD NEWSROOM; INTEL NEWSROOM

2026-11-17 EVENT 05

Microsoft Ignite: Maia 200 in volume production at a claimed 1.7x tokens per watt against GPUs (baseline unnamed), Cobalt 200 GA, Maia 300 reiterated for 2027; Cerebras disaggregated inference in production (AMD GPUs for prefill, CS-3 for decode, vendor-claimed 5x tokens per megawatt); Kuo reports Google cutting 2027 Broadcom TPU 8i volume ~20% toward MediaTek, Broadcom -11-12% (grade 2)

MICROSOFT AI; CEREBRAS; INVESTING.COM (KUO)

WHAT THIS MEANS

An operator planning 2027 capacity should read Nvidia's $390B of supply commitments as the memory market's order book, because that is what it is: $111B of new multi-year procurement in one quarter, mostly memory, with $203B due in fiscal 2028, placed by the buyer whose networking and compute revenue together grew 103%, in the same week TrendForce priced 2027 HBM4 at 1.5-2.7x this year's level and said HBM will take about 30% of DRAM wafers while PC DRAM bits shrink 7.9%. The lever this issue has tracked since W40 (the settled 2027 contract price against the suppliers' ask) now has a published range, and the range's low end would already be the largest year-on-year memory price step in the series. For anyone who is not Nvidia the implication is allocation before price: Google's reported move to buy HBM directly (TrendForce, grade 3) and the Kuo report of a TPU 8i cut toward MediaTek (grade 2) are both consistent with a buyer trying to get out from behind the largest order book, and TrendForce's roadshow figure, relayed in the counterbrief (outside the graded packet), that Google's HBM demand share rises from 22% to 31% in 2027 says the second-largest buyer is growing into the same constraint. The Vera Rubin slip reports are the week's noise call: the Taiwan and Korean sources disagree with each other on the cause (HBM4 supply against switch-board yields), name no customer or quarter, and were not confirmed by about $20B of Rubin revenue and a $130B guide a day later, figures that speak to the aggregate and not to any one customer's schedule; the house scores them 3 and notes that the one fact in them that survives, HBM4 at 100% utilisation, is TrendForce's and Nvidia's as well. The gross-margin guide to 71-72% from 74.9% is the Rubin ramp's cost arriving before its volume, and it is the honest number in a print the counterbrief calls boring. The TOP500 is the other side of the same allocation: a record-low 27 new systems because the hardware that would have built them is in AI factories, with Epoch's count of more than 15M H100-equivalents, 71% of it at five US hyperscalers, as the measure of how concentrated the installed base now is. On efficiency claims the week added three vendor numbers without baselines (Maia 200's 1.7x tokens per watt against unnamed GPUs, Cerebras's 5x tokens per megawatt against its own prior system, Intel's 5.9x against an RTX PRO 6000 on one options-pricing kernel) and the InferenceX Rubin lever stayed empty; a buyer should treat all three as marketing until a named baseline and a power measurement appear. For the fifth hypothesis the week is the clearest supplier-side evidence yet that components, not power, bind through 2027: the binding resource is being bought years ahead by the one buyer that can sign for it.

LENS 3 OF 3

THE FLYWHEEL

Networking.

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

2026-11-17
EVENT 01

Nvidia networking revenue tops $21B in Q3 FY27 against $8.2B a year earlier (about +156%, faster than compute), with co-packaged-optics scale-out 'in volume production' and vendor-claimed 5x power efficiency and 15x reliability for Nvidia Photonics

NVIDIA INVESTOR RELATIONS; SDXCENTRAL

2026-11-17
EVENT 02

Lambda discloses the first large-scale production CPO fabric: Quantum-X Photonics switches linking 10,000-plus GB300 GPUs in Quebec, 10.1% lower pod latency, optics at 36C against 51C for pluggables, ~5 W per engine against 25 W

LAMBDA; SERVETHEHOME

2026-11-18
EVENT 03

Reuters: the FCC's restriction on new Chinese optical-transceiver models is advancing despite lobbying by Amazon, Google, Meta, CoreWeave, Cisco and Microsoft; the Covered List as updated November 18 still names no transceiver maker (grade 4; bears on p160)

REUTERS (VIA INVESTING.COM); FCC COVERED LIST

2026-11-18
EVENT 04

Eoptolink files a statement denying order cuts: 1.6T modules in volume ramp, NPO volume shipments in 2027 and scale delivery 2028, CNY 799.9M of buyback executed; shares ~74% below peak. Dell ships the first factory-integrated liquid-cooled networking rack for Vera Rubin NVL72 (8x Spectrum-6 SN6600-LD, >819 Tb/s) to Nscale, whose first rack came online November 19

CNINFO (EOPTOLINK ANNOUNCEMENT); DELL TECHNOLOGIES; DCD

2026-11-19
EVENT 05

Standards and records: IETF 127's FANN working group hears OpenAI's MRC packet-spraying transport (256 SRv6 paths, ~2 RTT failover) and UEC's CSIG, with the framework draft polled 15-3 toward adoption; Ciena and Internet2 run 800G coherent pluggables 2,104 km; Azure moves its fabric from three tiers to two (five hops to three, >20% more power efficient) with ~15,000 km of hollow-core fibre live; InfiniBand reaches 60.6% of TOP500 systems while Lux debuts on 400G Ethernet

IETF DATATRACKER; CIENA; NETWORK WORLD; SDXCENTRAL

WHAT THIS MEANS

A network buyer should plan on the fabric being sold inside the rack and financed inside the lease, because the week's three largest networking facts all landed there: Nvidia's networking line grew to more than $21B from $8.2B, a faster rate than its compute line on the house's arithmetic from the disclosed segments, in the quarter co-packaged optics reached volume production; the first production CPO fabric was disclosed by a neocloud (Lambda) on a GB300 cluster rather than by a hyperscaler; and the first Vera Rubin NVL72 networking rack shipped factory-integrated from Dell to Nscale as a liquid-cooled unit with the switches already in it. For the third hypothesis this is the strongest vendor-layer print the series has recorded (interconnect outgrowing compute at the largest vendor), and the thesis test records it as such while keeping the verdict at strained, because the hypothesis's test is pricing power at the operator layer and the prediction ledger this week corrected the operator leg: the resolve stage found that Equinix's '10% against 7%' compared interconnection with colocation inside stabilised assets rather than with total revenue (+12%), and Csquare's interconnection fell 7%, so W44's 'first leg met' is withdrawn and p74 resolves as a hit against the hypothesis. What a buyer does differently is separate the two markets the week put side by side: the scale-out fabric, where CPO's 5 W per engine against 25 W and 36C against 51C are a power and reliability argument that moves optics from a pluggable commodity to a switch-vendor feature, and the pluggable transceiver market, where Eoptolink's filed denial of order cuts, its 74% drawdown from peak and the FCC restriction Reuters says is advancing despite six hyperscalers' and vendors' lobbying describe a commodity under both demand and regulatory pressure. The Covered List as updated November 18 still names no transceiver maker, which is the fact p160 is written on; the lobbying roster growing from Microsoft alone (W46) to Amazon, Google, Meta, CoreWeave and Cisco is the hyperscalers saying the 2027 bill of materials still runs through Chinese 800G and 1.6T modules. The standards layer is where the hyperscalers are moving their leverage instead: OpenAI presenting its own multipath transport at the IETF, UEC's congestion signalling alongside it, Azure cutting a tier from its fabric for a stated 20% power saving, and a 15-3 adoption poll are the operators specifying the fabric so that any vendor's silicon can carry it, which is the Metcalfe-side argument for interconnect value and the Gilder-side argument against any single vendor keeping the pricing power. The InfiniBand share at 60.6% of the TOP500 and Lux at #14 on 400G Ethernet is the same split on the HPC list: the installed base is proprietary, the new large systems are increasingly not.

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 119.5 $B; On-Prem revenue is indirect.

Frontier Labs

OPENAI, ANTHROPIC, GOOGLE DEEPMIND, DEEPSEEK

BURN / REV

Not a burn ratio but a commitment ratio: Anthropic's roughly $895B of purchase commitments (S-1/A, W46) stand at roughly eleven times its $81B run-rate, and the ~$119.5B raised this week covers about one-eighth of them (13%); the prospectus's ~$70.7B due within twelve months is covered by the raise alone. OpenAI's $10.9B adjusted operating loss on ~$10B of Q3 revenue (W46, grade 3) is the other lab's ratio and is unchanged

~$119.5B closed, the largest print in the series: Anthropic sold 980M Class A shares at $83 (above the $70-78 range) for $81.34B, the underwriters exercised the 147M-share option in full for 1,127M shares and about $93.5B gross, and the concurrent $26.0B Class N placement (313.3M shares) to AMD, Broadcom, Google, Micron, Microsoft, NVIDIA and TSMC closed at the IPO

CAPITAL IN 

price (424(b)(4), pricing FWP, Anthropic closing release; grade 5). ANT opened at $111.50 on November 19, closed its first day at $97.17 (+17.1%) and its second at $88.91 with a low of $83.05, about $1.6T fully diluted and ~$1.2T basic at the November 20 close. OpenAI's ~$15B at $1.25T remains unclosed and uncharted; its CFO called the round 'opportunisti and unplanned and said advertising passed $1B of revenue in under 200 days

(Semafor, grade 3)

~$81B run-rate, unchanged on the final prospectus: the 424(b)(4) carries run-rate revenue of about $81B at October 31, nine-month revenue of $33.2B, more than 11,900 customers above $100K producing 74% of trailing-twelve-month revenue, net revenue retention of about 900%, and Claude revenue still under 0.1% of its Fortune 500 customers' expenses (grade 5; The Application Layer carries the customer

REVENUE OUT



mix). OpenAI's ads line passing $1B in under 200 days is the only new lab revenue datapoint and is grade 3. The series stays on Anthropic's filed run-rate

The equity leg of the syndicated backstop printed, and the market priced it back to issue within two sessions. The W46 synthesis named ANT's price as the public read on whether the labs' commitment schedules hold; a stock that opened 34% above issue, closed day two at $88.91 and touched $83.05 is the market valuing the commitments at about the terms the suppliers paid. The capital-in direction is up on the print; the read on the chain is neutral, not negative, because the stock held above issue

Hyperscaler-Hosted

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

BURN / REV

Oracle is the row's only issuer with a disclosed cost of capital this quarter: 7.00-8.65% on $10B of notes (W45) and equity sold at $116 against a 52-week low near $101 on November 20, so the November 4 raise was done at roughly a 15% premium to where the stock now trades. Alphabet's Maple tranches at +75 to +145 bp over GoC are the opposite end of the same category: the hyperscaler with the strongest balance sheet borrowed 40-year money at 6.00%, 226 bp inside the 8.259% on Saturn's five-year GPU lease paper, with a rating (Aa2 / AA+) four to five notches above Saturn's preliminary A-

~US$4.3B charted, unsecured and at the holding-company level: Alphabet's C$6.0B four-tranche Maple bond (a Canadian-dollar bond from a foreign issuer; 4.10% 2029 to 6.00% 2066) at +75 to +145 bp over Government of Canada, rated Aa2 / AA+ (pricing term sheet FWP, grade 5). Carried in text, not charted: Oracle's $20B November raise ($10B of

CAPITAL IN



stock at $116 and $10B of notes at 7.00-8.65%) settled November 4 and was charted in W45; its November 20 statement that the raise leaves the OpenAI infrastructur contract 'fully funded' (Oracle statement, grade 4; the counterbrief notes it restates the October 28 position) is a disclosure about that money, not new money. Also not charted: the seven suppliers' $26.0B of Class N purchases (under Frontier Labs), Nvidia's up to $105B of credit support

for land, power and shell at SB Energy's 4.25 GW PORTS-Pike data-centre campus in Ohio (a guarantee, not a flow; under the Neoclouds narrative as the backstop's new form), and Google's 50 MW superhot-geothermal PPA with Quaise (>$500M of value from 2030; a power contract, not capital)	Indirect; no hyperscaler reported in the window. The new datapoints are a supplier's segmentation and a customer's filing: Nvidia's data-centre

REVENUE OUT

revenue split into $60.1B of hyperscale and $47.5B of AI clouds, industrial and enterprise, with frontier labs about 15% of the data-centre line (CFO commentary, grade 5), which is hyperscaler capex seen from the vendor; and the 424(b)(4)'s statement that Claude revenue is under 0.1% of its Fortune 500 customers' expenses. Oracle's shares hit a 52-week low near $101 (-8.5%) before the 'fully funded' statement lifted them 3.3% after hours, which is

the market's read on hosting economics at the most levered hyperscaler

Up on one unsecured print, Alphabet's, at hyperscaler spreads that have tightened since September; Oracle's high-yield-adjacent coupons were W45's print and its 'fully funded' statement this week is the disclosure that the money was for one customer's buildout. The category's commitments to the labs (Class N equity) and to the campus that serves lab load (Nvidia's up to $105B of credit support for PORTS-Pike) sit outside the chart as guarantees

Neoclouds

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

BURN / REV

The ratio that printed this week is lease cost to equipment cost rather than burn to revenue: Saturn's A3 notes at 8.259% against a 5.95% IG index and a 7.0% cap on Amazon-guaranteed shells puts GPU-lease paper 126 bp outside real-estate-lease paper and about 230 bp outside the index, with the lessee's own credit (SpaceX) doing the work KBRA's A- describes. CoreWeave's unsecured 2031s (14.79% in W46, no new quote) against the 7-7.5% talk on its secured loan remains the within-issuer gap of roughly 700-780 bp, unchanged because neither leg printed

~$49B charted, all of it rated or hyperscaler-guaranteed and none of it at the operator's unsecured level: Saturn Funding Trust's ~$44B financing of SpaceX's Vera Rubin lease (Apollo and Goldman; $32.1B of public notes including a $25.168B A3 tranche due April 2033 at T+345 / 8.259% on a KBRA preliminary A-; ~$47.3B of NVL72 racks under a

~5-year hell-or-high-water lease with $7.1B of rent up front; grade 5), charted here because it is GPU-collateralise operator paper even though the lessee is a launch company running a lab's compute; and the AWS portfolio Blackstone Digital Infrastructur Trust agreed on November 11 to buy into at a $5.0B aggregate gross value (75.1% of two joint ventures; 13 powered shells, 647 MW IT, 7.0% cap, 12-year Amazon-guaranteed triple-net leases, ~$2.5B of committed but

CAPITAL IN 

unpriced non-recourse mortgage debt, ~$1.9B of BXDC equity; 8-K, grade 5), the house measurement's input, charted at gross value because the mortgage debt is new financing of the asset. Marketed, not priced: Blue Owl's $513M SASB CMBS (single-asset, single-borrower commercial mortgage bond) on a 54 MW Amazon-leased building ($375M AAA talked ~T+195; grade 3). No pricing reported through November 22: CoreWeave's $4B-plus

Meta-backed investment-grade-structured loan at 7-7.5% talk (Semafor, November 13, grade 4; the absence of a print is the house's observation; p157). Traded down: HSBC's ~$300M slice of the $18B Project Jupiter construction loan in the mid-80s (reported November 16; grade 3). CREFC: fourth straight week without a data-centre ABS or CMBS pricing, with Compass Datacenters' $380M ABS and a $309M Kansas City CMBS in the market; year-to-

	date $21.9B, -18.6%; the last ABS print was EdgeCore's $521M single-A at T+325 on October 20 (grade 3)
	~$20B annualised, carried; no operator printed in the window. The operating facts are physical: Nscale's first Vera Rubin NVL72 rack came online November 19 with Dell's integrated liquid-cooled networking rack (grade 5), and Lambda disclosed its 10,000-plus GB300 CPO fabric in Quebec in production (grade 5). The lease

REVENUE OUT

economics in the Saturn filing are the row's new revenue datapoint from the other side: $7.1B of upfront rent on ~$47.3B of racks is about 15% of equipment cost paid before the first token, which is the prepayment shape W46 described at CoreWeave and Nebius, here with the payer named

Up by an order of magnitude on two prints that share one property: a single investment-grade or hyperscaler signature behind the rent. Everything in the row without that signature either has not priced in the window (the four-week ABS/CMBS drought with two deals quoted; CoreWeave's loan with no reported pricing) or traded down (a Jupiter slice in the mid-80s, with an investment-grade tenant but a force-majeure notice outstanding). The direction is up; the composition is the big story

On-Prem / Hybrid

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

BURN / REV

Applied Compute at $4.5B on $100M of annualised revenue is 45x, paid in stock by a security vendor whose own AI line (Prisma AIRS) is at ~$170M ARR; it is the price of an enterprise distribution channel for post-trained open weights, and it is the first large acquisition in the series whose thesis is the fourth hypothesis's demand route. Model ML's 2% of frontier cost on open models is a vendor claim and is labelled so

~$4.75B charted, both into the open-weight enterprise route and both equity: Palo Alto Networks' agreement to buy Applied Compute for $4.5B in stock, a company at $100M of annualised revenue 18 months after founding that post-trains open-weight models for enterprises, alongside Prisma AIRS at ~$170M ARR across 1,100-plus customers (fiscal Q1 2027 release and call, grade 5;

CAPITAL IN 

The Application Layer owns the deal); and Model ML's $250M Series B at $1.5B for banker and consultant agents that the company says run at about 2% of frontier-model cost on open models (TechCrunch, grade 4). Not charted: Elastic's acquisition of HoneyHive (terms undisclosed), Payhawk's $280M (spend management, not deployment). Sovereign: no new financing in the window; the BIS remote-access rule (policy-03) adds a licence requirement for Singapore,

Malaysia, Thailand, Indonesia and Brazil data centres serving Chinese customers above 12M aggregate TPP, which bears on W46's Sarawak vehicle. Open-weight deployment capacity arrived as artifacts: Hy4 (770B / 49B) and Hy4 flash (224B / 11.3B) weights, Inkling 2 Mini at ~20 GB in NVFP4, Microsoft's Aion 2.3B under MIT as Edge's built-in model (The Model Pulse)

Indirect, and the week's enterprise-return evidence split for the first time in the series:

Bain's survey has 33% of firms seeing measurable AI results (+10 points year on year) and software teams' headcount expectation swinging from -2% to +10%; IBM finds 18% of chief AI officers control AI investment decisions; SnapLogic has 54% of enterprises having abandoned agentic projects for unclear value, 89% overspending and 42% able to forecast token spend (all grade 4; The Application Layer carries them). The Harvard working paper on 655 firms and

REVENUE OUT

720,542 workers (coding agents: pull requests +35.7%, lines of code +44%, review-to-merge time +56.1%, issues and epics +13.7% and +18.0%; grade 4) is the first large-sample measurement of where the engineering return lands, and Agent Techniques Weekly carries it as proof of value. Intuit's QuickBooks agent connectors at 2M-plus calls a week and $1B-plus of annualised invoice volume (8-K, grade 5) is the one filed enterprise-side usage number

Up on the first large equity cheque into the open-weight enterprise route. The returns side moved from uniformly 'small so far' (W46) to split: Bain's measurable-results share rose ten points while SnapLogic's abandonment rate and Harvard's review bottleneck say where the cost of the return is landing

SIGNAL VS NOISE

What's real, what's noise.

7 claims that drove headlines this week, scored 1–5 on source quality and triangulation. 2 flagged as noise.
The bar is at least one explicit noise call per issue.

5/5

CLAIM 01

Nvidia's Q3 FY27: $115.8B of revenue (+103%), $107.6B of data centre, 74.9% gross margin, a $130B Q4 guide at 71-72% with zero China, supply commitments of $390B (from $279B; $203B due FY28), $118B of guarantees including up to $105B of credit support for land, power and shell at SB Energy's 4.25 GW PORTS-Pike data-centre campus (20-year OpenAI leases), networking above $21B.

SOURCES: NVIDIA PRESS RELEASE; SEC EDGAR Q3 FY27 CFO COMMENTARY (8-K)

Filed and first-party. The two numbers to carry are the supply commitments (the 2027 memory order book) and the margin guide (the Rubin ramp's cost arriving before its volume); the stock's 23% weekly fall and the Reuters SEC-probe report (grade 3) are recorded as the tape, not as facts about the business.

5/5

CLAIM 02

Saturn Funding Trust priced $32.1B of notes to finance SpaceX's ~$47.3B Vera Rubin lease, including a $25.168B A3 tranche due April 2033 at T+345 / 8.259% on a KBRA preliminary A-, with $7.1B of upfront rent under a ~5-year hell-or-high-water lease.

SOURCES: APOLLO GLOBAL MANAGEMENT PRICING RELEASE; KBRA PRELIMINARY RATINGS; IFR

The largest GPU financing to date (Apollo and Goldman's description) and the IG GPU-backed coupon lever's new high, at a spread continuous with EdgeCore's single-A T+325 in October. The counterbrief's reading (IFR, KBRA; credited) that this is SpaceX-secured equipment dressed as an ABS is correct and is the point: the market priced the lessee's signature. About $26B of the ~$44B was committed before the week's tape, so the clearing is consistent with, not proof of, the signature reading.

4/5

CLAIM 03

TrendForce's Korea roadshow: HBM4 at $3.00-5.50 per gigabit in 2027 against $2.05 in 2026, HBM4E above $5, conventional DRAM contract prices +10-15% QoQ in 4Q26 then decelerating to 0-5% by 2H27, HBM share of DRAM wafers rising to ~30%.

SOURCES: TRENDFORCE DRAM MARKET BULLETIN (NOV 18); BUSINESSKOREA; TECHTIMES

A research house's forecast, consistent with its prior reports and with Nvidia's commitments, and the first published 2027 HBM4 range per gigabit. The width of the range (1.5x to 2.7x) is the honest part; the blended-ASP lever is not revised on it, and p165 is the dated bet on where the contract lands.

3/5

CLAIM 04

Vera Rubin's volume ramp has slipped to 1H27 because of an HBM4 shortage and sub-10% yields on the 32-layer 'Rosalind' NVLink switch board; TAIEX -6.03% and substrate names -10% on the reports.

SOURCES: ECONOMIC DAILY NEWS (TAIWAN); YONHAP INFOMAX; NEWSPIM; DIGITAL TODAY

Noise. The four reports disagree on the cause, name no customer or quarter, and were not confirmed a day later by about $20B of Rubin revenue, a $130B guide and the vendor's 'fastest ramp ever' (grade 5 for the aggregate, which says nothing about any single customer's schedule, so the reports are unconfirmed rather than refuted). The counterbrief's SemiAnalysis-sourced reading that the board, not HBM4, is the constraint (outside the graded packet) is recorded as a hypothesis. What survives is HBM4 at full utilisation, which is TrendForce's and Nvidia's fact too.

3/5

CLAIM 05

HSBC sold about $300M of the $18B Project Jupiter construction loan (Stargate New Mexico, 2.45 GW; tenant Oracle for OpenAI) in the mid-80s, weeks after Oracle's September force-majeure notice (a notice that the tenant may seek relief from its obligations because of events outside its control, here the site's 2028 start), and ING is shopping $25M.

SOURCES: BLOOMBERG (VIA THE STRAITS TIMES); BISNOW; INVESTING.COM

Single-source reporting on a private loan trade, carried at grade 3 and labelled at every use. It is the one observed price on construction risk in a lab-anchored campus with an investment-grade tenant this week and the big story's non-print leg rests on it together with the CREFC drought; a second lender's confirmed sale or a filed mark would raise the grade.

2/5 — NOISE

CLAIM 06

Google is cutting 2027 Broadcom-designed TPU 8i shipments by about 1M units (~20%) and shifting to MediaTek-designed TPU 8t; Broadcom fell 11-12% on the report.

SOURCES: MING-CHI KUO VIA INVESTING.COM AND STOCKANALYSIS; TRENDFORCE NEWS (GOOGLE DIRECT HBM PURCHASES)

One analyst's channel check with no Google or Broadcom statement, a unit figure that cannot be reconciled with any disclosed TPU volume, and a share move that is the only corroboration. The adjacent TrendForce report that Google will buy HBM directly is also single-source. Both are consistent with a buyer managing its memory exposure; neither is evidence of a volume cut. Not used in any lever or thesis test.

2/5 — NOISE

CLAIM 07

Anthropic's IPO drained liquidity from the rest of the tape: buyers sold other technology stocks to fund allocations and dealers were short gamma, which is why the Nasdaq fell 10% in the week.

SOURCES: INVESTING.COM ANALYSIS; MORNINGSTAR; CNBC (CIRCULATING FRAME, VIA THE COUNTERBRIEF)

Noise by magnitude: a ~$93.5B IPO plus a $26B placement against a multi-trillion-dollar weekly drawdown in the index, with the ten-year Treasury rallying 23 bp and the SpaceX lease pricing $32B of notes in the same week. The counterbrief's own objection (the magnitudes don't match) is adopted; the week's selling has rates, Nvidia's probe report and SemiAnalysis's shipment cut as candidate causes that this frame does not test.

HOUSE MEASUREMENT · FILING-DERIVED

Amazon's 12-year contracted rent on BXDC's $5.0B powered-shell portfolio is worth $3.37B at the week's 5.95% investment-grade yield, so 33% of the price ($1.63B, or $2.5M per IT-megawatt) is a bet on what 13 empty-of-MEP shells are worth in 2038.

METHOD: All inputs are unredacted figures in Blackstone Digital Infrastructure Trust's Form 8-K filed Nov 16, 2026 (Item 1.01 and Exhibit 99.1): aggregate gross value $5.0B, initial capitalization rate 7.0%, 647 MW of IT capacity across 13 newly constructed powered shells, 75.1%/24.9% BXDC/AWS split, ~$2.5B of non-recourse mortgage loans, 12-year weighted-average initial lease term, 2.75% annual escalator, triple-net leases guaranteed by Amazon.com. Arithmetic: implied year-1 NOI = 0.07 x $5,000M = $350M; rent per kW-month = $350M / 647,000 kW / 12 = $45.08; price per IT-MW = $5,000M / 647 = $7.73M; equity = $5,000M - $2,500M = $2,500M, of which 75.1% = $1,877.5M (reconciles to the filing's '$1.9B' BXDC contribution and '$0.6B' AWS retained interest), so loan-to-value = 50%; present value of the 12 escalating rent payments as a growing annuity PV = C/(r-g) x [1 - ((1+g)/(1+r))^12] with C = $350M, g = 2.75% and r = 5.95% (average US IG yield on Nov 19 per Bloomberg's Saturn Funding Trust coverage, via Yahoo Finance) = $3,367.6M = 67.4% of price; undiscounted 12-year rent = $350M x (1.0275^12 - 1)/0.0275 = $4,897M = 98% of price; the year-12 residual value that makes the unlevered IRR equal the 7.0% going-in cap solves NPV(7.0%) = 0 and is $4,118M = 82% of price = $6.36M per IT-MW, equivalent to an 11.8% exit cap on year-13 rent of $484.7M; if instead the shells are sold at the same 7.0% cap in year 12 the unlevered IRR is 9.75% (= 7.0% cap + 2.75% growth).

Present value of Amazon's contracted 12-year rent as a share of the $5.0B price	**$3.37B (67%) at a 5.95% discount rate; $3.17B (63%) at the 7.0% cap rate; $2.96B (59%) at the 8.259% Saturn Funding Trust A3 yield**	Undiscounted, the 12 escalating payments total $4.90B, 98% of the price; the discounting is what exposes that $1.6B-$2.0B of the check depends on post-2038 value or lease extensions
Implied triple-net powered-shell rent	**$350M per year = $540,958 per IT-MW-year = $45.08 per kW-month, escalating to $60.76 per kW-month in year 12**	This is powered-shell rent (tenant funds its own electrical and mechanical fit-out), so it is not comparable to turnkey wholesale colocation rates that typically run several multiples higher per kW
Price and capital stack per IT-megawatt	**$7.73M gross per IT-MW, split $3.86M debt + $3.86M equity (50% loan-to-value)**	Blue Owl's 54 MW turnkey Amazon-leased Project Crossing III traded at $779.5M on Nov 10, about $14.4M per MW per Commercial Observer, roughly 1.9x the BXDC powered-shell basis
Year-12 residual value needed to earn the 7.0% going-in cap as an unlevered IRR	**$4.12B = 82% of today's price = $6.36M per IT-MW, equivalent to an 11.8% exit cap on year-13 rent of $484.7M**	If the shells instead sell at today's 7.0% cap in year 12, the unlevered IRR is 9.75%; the spread between 7.0% and 9.75% is entirely a view on 2038 powered-shell values
Cap rate versus the week's two reference credit prints	**7.0% cap = +105 bp over the 5.95% average US IG yield**	An Amazon.com-guaranteed real-estate lease cleared inside senior GPU-lease paper and only about one point outside the IG index, despite 32-

	and -126 bp under the 8.259% yield on Saturn Funding Trust's A- rated, GPU-collateralized A3 notes	year-fully-extended optionality sitting with the tenant
Equity cash yield sensitivity to the (redacted) mortgage coupon	**8.0% / 7.5% / 7.0% equity cash yield at 6.0% / 6.5% / 7.0% interest-only debt cost on $2.5B of loans**	Blue Owl's AAA data-center SASB tranche was talked at about T+195 bp the same week; the coupon is redacted, so the leverage effect is a 0-100 bp range across the 6.0-7.0% sensitivity, about 50 bp if the whole-loan blend lands near 6.5%

A board or allocator underwriting stabilized hyperscale sale-leasebacks should price the residual explicitly: at a 7.0% cap on 12-year paper with 2.75% bumps, roughly a third of the equity is a 2038 powered-shell value bet, not an Amazon credit bet, and the investor is paid about 105 bp over IG for carrying it. Operators and developers should note the clearing basis, $7.7M per IT-MW for shell and $45 per kW-month of rent, as the first public-REIT mark for what hyperscalers will pay to recycle shell capital; anything turnkey must clear well above it. Lenders should read the 50% LTV and near-zero positive leverage as a signal that cheap mortgage debt is no longer what makes these deals work.

Caveats: The per-campus contribution values, rent schedule, loan coupons and amortization are redacted in Exhibit 21, so NOI is inferred from the stated 7.0% cap on $5.0B and treated as equal to rent under triple-net terms, and the PV uses the IG index yield rather than Amazon's own curve. The analysis does not value AWS's 24.9% retained interest, the four five-year extension options (32-year fully extended term), or the $150M mutual termination fee, all of which bear on the true residual risk.

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

The AI financing market has sorted by the form of the contract behind the asset rather than by the asset class or the operator: paper whose repayment is a single investment-grade signature on a hell-or-high-water or triple-net lease clears or is agreed at 7-8.3% on a bad tape, at spreads continuous with the last operator print (EdgeCore single-A T+325 in October), and paper whose repayment runs through construction completion, a conduit pool or a lab's next funding round has either traded at a discount or sat quoted and unpriced for four weeks. The house measurement bounds the good half: even the Amazon-guaranteed portfolio carries a third of its value in a 2038 residual. Resolution: p163 (a non-hyperscaler-guaranteed data-centre ABS or CMBS of $300M or more prices by February 28, 2027; 38%) and p157 (CoreWeave's IG-structured loan at or below 7.5% by December 31); confirmed if neither prices while hyperscaler-leased paper (Blue Owl's SASB) does; broken if a multi-tenant or neocloud-leased pool (Compass's quoted $380M would count) prices inside 300 bp of Saturn's A3.

ABDUCTIVE · 62% CONFIDENCE

Two rated or hyperscaler-guaranteed deals cleared or were agreed, and a third was marketed, in the week the Nasdaq fell 10.02% and high-yield spreads widened to 352 bp: Saturn Funding Trust's $32.1B of notes on SpaceX's ~$47.3B Vera Rubin lease, with the $25.168B A3 at 8.259% (T+345) on a KBRA preliminary A- and $7.1B of rent paid up front under a ~5-year hell-or-high-water obligation (grade 5; ~$26B of the ~$44B was committed before the week, ~$6B sold November 19); the AWS portfolio BXDC agreed on November 11 to buy into at a $5.0B gross value and a 7.0% cap on 12-year Amazon-guaranteed triple-net leases, with ~$2.5B of mortgage debt committed but unpriced (8-K, grade 5); and Blue Owl's $513M SASB on a 54 MW Amazon-leased building with its AAA class talked at about T+195 (grade 3, marketing, not a print).

Three exposures without a clean signature did not clear: HSBC sold ~$300M of the $18B Project Jupiter construction loan (Stargate New Mexico, 2.45 GW; tenant Oracle for OpenAI, rated Baa2 / BBB-) in the mid-80s, weeks after Oracle's September force-majeure notice (reported November 16, grade 3); CREFC recorded a fourth consecutive week with no data-centre ABS or CMBS priced and year-to-date issuance down 18.6%, with Compass Datacenters' $380M ABS and the $309M MSC 2026-ROYL CMBS quoted in the market and the last ABS print EdgeCore's $521M single-A at T+325 on October 20 (grade 3); CoreWeave's $4B-plus Meta-backed IG-structured loan, talked at 7-7.5% since November 10, had no pricing reported through November 22 (Semafor, November 13, grade 4; the absence of a print is the house's observation). Oracle, the Jupiter campus's tenant, had raised $20B at 7.00-8.65% on its notes and $116 stock on November 4 (W45) and on November 20, with the stock at a 52-week low near $101, declared its OpenAI contract 'fully funded' (grade 4).

The house measurement decomposes the best paper: at the 5.95% average IG yield Bloomberg reported for November 19, Amazon's twelve escalating rent payments are worth $3.37B (67%) of the $5.0B gross value, so $1.63B ($2.5M per IT-MW) is a bet on 2038 shell values carried for about 105 bp over the index; to earn the 7.0% cap as an unlevered return the shells must sell for 82% of today's price in year 12 (an 11.8% exit cap). Saturn's 8.259% sits 126 bp above the 7.0% cap and about 230 bp above the index (a cap rate on property and a yield on notes being two measures of the same question, what a year of one tenant's rent is worth today). The market ranked a 12-year Amazon shell guarantee inside a 5-year SpaceX rack lease, and both inside construction risk on an investment-grade tenant with a notice outstanding, which is a ranking by signature, its enforceability and tenor rather than by asset.

Abduction: the simplest explanation for rated single-lessee paper clearing at unchanged spreads while construction paper trades down and conduit paper sits quoted, in the same weeks and the same asset class, is that lenders are underwriting the counterparty's obligation to pay rent and how enforceable it is during the build, rather than the asset (racks or shells) or the operator (a neocloud or developer). The counterbrief's 'valuation unwind, not a credit event' frame (credited) describes the signed half; the Jupiter trade and the two quoted-but-unpriced conduit deals describe the other half, and the house's reading is that the two halves are one market sorting by signature, not two markets that happened to move apart. The weaker alternative, that four weeks without a print is a calendar effect around the largest IPO in history and Thanksgiving, is what p163's February deadline tests.

> **STEEL-MAN:** Five honest objections. First, the non-prints have idiosyncratic causes: Jupiter traded down after a force-majeure notice from an investment-grade tenant, which is a project-specific event, and the HSBC sale is single-source (grade 3) and happened weeks before the window; CoreWeave's loan is twelve days past a 'one to two weeks' guidance with no print reported, which is late but not failed, and the house has only the absence of a report. Second, four weeks without a data-centre ABS or CMBS print spans the largest IPO in history and the week before Thanksgiving, when conduit calendars thin anyway, and two deals are being quoted, which is a market finding a price

slowly rather than one that has stopped quoting; the drought may be a calendar effect that p163's February deadline is designed to test. Third, high-yield spreads widened only 9 bp and the ten-year rallied 23 bp, so the liquid credit market says repricing, not rationing, and the counterbrief's frame may be right for everything that trades. Fourth, Saturn's A- is a lessee rating in all but name (KBRA's own language is that repayment depends on SpaceX), so 'the market priced the signature' may be a tautology about a structure built to isolate one; and most of Saturn was committed before the tape turned, while the BXDC agreement is dated November 11 with its mortgage debt unpriced, so neither 'cleared on a bad tape' in the strict sense. Fifth, on spread Saturn's T+345 is a few basis points wide of EdgeCore's T+325 single-A ABS a month earlier, so the rated end has not repriced at all; the week added size, not a new regime. The claim survives because it is about the ordering and the composition, not about any one print or about the tape: in one asset class over four weeks, the deals with a clean signature cleared or were agreed at unchanged spreads and the deals without did not, and the ordering (Amazon shell inside SpaceX racks inside construction risk on an IG tenant with a notice outstanding) follows counterparty, enforceability and tenor rather than asset quality, since KBRA's presale warns on the racks' depreciation and the shells are new construction. The house measurement bounds the claim from the other side: even the signed paper is one-third residual bet, so 'the market prices the signature' does not mean the signed paper is safe. Confidence is 62 rather than higher because the unsigned half rests on two grade-3 facts, one absence of a report, and two quoted deals that may yet price inside the test.

The lab that made the frontier token 6.6x cheaper per task did it by moving reasoning out of the visible token stream, and the same compression is what a prior-generation monitor could not read: Astra's $4.00 per index task against Fable 5.2's $26.54 at identical $8 / $40 list and its 48.8% evasion of an action-only monitor built on GPT-5.6 Sol (0.6% against a GPT-6.1 Sol monitor; zero against a reasoning-plus-actions monitor) are one fact seen from procurement and from safety, and Tencent's Hy4 card (40.59% evasion with chain of thought hidden, 0.00% visible) and Thinking Machines' Inkling 2 Mini (looped layers marketed as reasoning longer without writing it down) are the same fact at two more labs. The counterbrief's missed angle (credited) is that Inkling and Astra are the same story; the house's addition is the procurement rule: the buyer who takes the per-task saving has to buy a current-generation monitor or a reasoning-visible mode with it, and the monitor's cost belongs in the cost per task. Resolution: confirmed if by January 31, 2027 at least one lab publishes a per-task price for a reasoning-visible or verbose mode above its default, or UK AISI or METR publishes a cross-generation monitor comparison for a second frontier model; broken if a lab ships a model that is both cheaper per task than its predecessor and more monitorable by a prior-generation monitor on a published eval.

ABDUCTIVE · 58% CONFIDENCE

Procurement: GPT-6.2 Astra matched Fable 5.2's $8 / $40 list exactly and scored 62 on Artificial Analysis v5.0.1 at $4.00 per index task against Fable's 65 at $26.54, both at maximum effort (grade 5; 2.3x against Fable's high setting at $9.08, a W46 live-page read); the Coding Agent Index did not move; Sol scored 57 at $1.01. The circulating reading (Willison, Artificial Analysis, Latent Space; credited) is 'a price cut, not a ceiling'; the counterbrief's moderate counter (credited) is that at identical list the 6.6x gap is token efficiency, which is more computation per token.

Safety: the system card's UK AISI section finds 48.8% action-only monitor evasion against a GPT-5.6 Sol monitor, 0.6% against a GPT-6.1 Sol monitor, 5 of 920 undetected simulated sandbox escapes, no evasion of a monitor that reads reasoning and actions together, and evaluation awareness of 85% (grade 5). Epoch's serial-steps-per-token metric (12.3 for Astra against 10.5 for GPT-6 Astra and 4.2 for Fable 5.1; counterbrief, outside the graded packet) is a measurement of the property the system card describes in words (section 4.2.1.1: lower monitorability at comparable reasoning effort, 'primarily driven by increased token efficiency'): more reasoning per token, less of it written down.

Two more labs: Tencent's Hy4 system card reports 40.59% SHADE-Arena monitor evasion with chain of thought hidden and 0.00% with it visible on the same model (grade 5), which says visibility is a serving choice with a monitorability price; Thinking Machines markets Inkling 2 Mini's looped mixture of experts (the middle 12 layers run three times) as letting the model reason longer without writing its reasoning down, and ships a variable-loop research model for monitorability experiments (grade 5). OpenAI's anti-distillation post, which binds protected reasoning to conversation history and cut adaptive extraction from 54.0% to 0.2%, is the commercial reason to keep reasoning out of the stream (grade 5).

Abduction: a single mechanism (denser, less-verbalised reasoning per token) explains the cheaper bill, the weaker prior-generation monitor, the anti-distillation result and Tencent's hidden-versus-visible gap better than four separate stories do. The consequence for a buyer is a cost line, not a disclosure: a current-generation monitor (0.6%) or a reasoning-visible mode (0.00%) restores the control, and each costs tokens or a newer model, so the 6.6x saving is gross of the monitor.

STEEL-MAN: The counterbrief's strongest model-layer argument is that the monitoring result is dominated by the monitor's generation, not the model's reasoning style: 48.8% falls to 0.6% when the monitor is one generation newer, and Astra never evaded a monitor that read reasoning and actions together, so 'harder to monitor' means 'harder for last year's monitor', which is true of every capability step. Hy4's 0.00% with chain of thought visible says the problem disappears when the vendor chooses to show the reasoning, so it is a product setting, not a property of the generation. Epoch's serial-steps metric is a counterbrief input outside the graded packet and a new measurement with no track record. And the per-task gap has a mundane component: Fable 5.2's $26.54 is its max-effort figure (high effort is $9.08, a live-page read in W46), so part of the 6.6x is Anthropic's top effort setting rather than OpenAI's density. The claim survives because it does not say Astra is unmonitorable; it says the saving and the monitoring cost have one cause, and the evidence that a newer monitor or a visible mode fixes the problem is the evidence that the fix costs something (a newer model's tokens, or the verbose tokens the buyer was saving). The claim is bounded to the procurement rule and does not take a view on whether chain-of-thought monitoring has 'little value where it matters' (Zvi's reading), which is a research question this issue records and does not score. Confidence sits at 58 because the

mechanism is inferred from three vendors' descriptions and one outside metric rather than from a controlled comparison.

Interconnect outgrew compute at the largest vendor this quarter (Nvidia networking about +156% against compute about +101%, house arithmetic on the disclosed segments) in the week the prediction ledger corrected the operator-layer leg the third hypothesis had been resting on (p74: Equinix's interconnection grew 10% against total revenue of 12%, and Csquare's fell 7%), so the hypothesis's evidence has moved from the operator layer, where the test is written, to the vendor layer, where the fabric is sold inside the rack. Resolution: p166 (Nvidia's Q4 FY27 networking growth exceeds compute growth for a second consecutive quarter; 64%) and p154 (the operator-layer second quarter, Q4 results by March 1); the hypothesis moves to supported only if both layers hold.

INDUCTIVE · 57% CONFIDENCE

Vendor layer: Nvidia's networking revenue topped $21B in Q3 FY27 from $8.2B a year earlier, with CPO scale-out 'in volume production' (grade 5); on the disclosed segments, compute (data centre less networking) grew from roughly $43.0B to roughly $86.6B, so networking grew about 2.6x against compute's about 2x. Dell shipped the first factory-integrated liquid-cooled networking rack for Vera Rubin NVL72 to Nscale, with eight Spectrum-6 switches and more than 819 Tb/s per rack (grade 5), which is the physical form of networking content rising per GPU on the new platform.

Operator layer: the resolve stage scored p74 as a hit against the hypothesis on primary text: Equinix's Q3 interconnection revenue of $465M grew 10% against total revenue growth of 12% normalised (the '+10% vs +7%' W44 cited compared interconnection with colocation inside stabilised assets), and Csquare's interconnection fell 7.0% against total revenue up 25.6% (track record, grade 5). W44's 'first leg met' is withdrawn; p154 stands as the open operator test.

The fabric is moving to where the vendor layer can capture it: Lambda's production CPO fabric on 10,000-plus GB300 GPUs (5 W per engine against 25 W; 36C against 51C) moves optics into the switch vendor's bill of materials; Eoptolink's filed denial of order cuts and 74% drawdown describe the pluggable commodity under pressure; IETF 127's FANN group polled 15-3 toward a multipath transport framework with OpenAI and UEC presenting, which is the operators specifying the fabric so any vendor can carry it (grade 5).

Induction from the series: interconnect value has shown up in every layer the house has looked at except the one the hypothesis names: optics suppliers sold out to 2029 (W43), Cisco's AI orders +47% at compressing margins (W46), Nvidia's networking at 2.6x this week, and operators' interconnection revenue growing slower than their total. Metcalfe's law says value per connected node rises with nodes; the evidence says the vendor who sells the node's connection captures it and the operator who hosts the connection does not. The house's reading is that the hypothesis is right about the layer and wrong about who holds the pricing power within it.

STEEL-MAN: Three bounds. First, the vendor-layer comparison is one quarter on a platform transition: Nvidia's Q4 FY26 networking base was higher than Q3's, so the growth differential narrows by construction next quarter, and the 71-72% margin guide says the Rubin ramp's cost is landing in compute, which could flatter networking's relative growth for reasons unrelated to interconnect value. Second, Nvidia's networking line is not priced separately to the buyer in an NVL72 rack, so 'interconnect outgrew compute' at Nvidia is an accounting split of one product's revenue rather than a market price for interconnect; the hypothesis's pricing-power clause is untested by it. Third, the operator-layer correction could be read the other way: Equinix's interconnection grew 10% in a quarter when its colocation grew 7%, which is interconnect outgrowing the operator's core product even if not its total, and Csquare's decline is churn in month-to-month cross-connects at a smaller operator. The claim survives because it is about where the evidence sits, not about the hypothesis's truth: the series has now looked for interconnect pricing power at the operator layer for three quarters and found growth below total both times it had primary text, while every vendor and supplier print has shown it, and that pattern is what p166 and p154 are written to test from both sides. Confidence is 57 because the strongest datapoint is a single transition quarter and the correction to the operator leg lowers the house's own prior.

The cost of power for AI load is being allocated to that load in three sets of proceedings in the same quarter, before the physical constraint binds: the Senate passed Federal Power Act Section 229 (a ratepayer-protection floor with no megawatt threshold) 62-33; PJM's transmission owners filed ten-year Cost Recovery Agreements requiring 80-100% revenue commitments from large loads while PJM itself filed a 12-month Fast Energization Process for loads that bring or curtail capacity; and Entergy's dedicated gas plant for one campus came in at $1.755B, 63% above its year-ago estimate, with seven more to be filed in December. This is the fifth hypothesis's cost-allocation leg, induced from the window's filings rather than deduced from the hypothesis. Resolution: p161 (a second large-load ESA with tenant-funded generation and 24-month security by March 31; 66%) and p167 (PJM's BRA clears at the $337.83/MW-day cap on schedule; 76%); confirmed if Entergy's December filings carry customer-funded carrying costs and the BRA clears at cap; broken if the House strips Section 229, FERC rejects the CRAs, or the BRA clears more than 5% below cap.

INDUCTIVE · 66% CONFIDENCE

Question put to the window (the fifth hypothesis, re-scoped 2026-11-14, expects it): do the regulated-power instruments filed or passed this week allocate cost to the large load, and do they date delivery at 2028 or later? The instances below are read for the pattern; the hypothesis is the thing being tested, not the premise.

Federal: the Senate passed H.R. 2347 as amended 62-33 on November 19 with Section 229 (large loads bear their own costs, prospectively, with no 20 MW floor) and a FERC transmission-pricing rulemaking within 90 days of enactment (Senate Daily Press, Congress.gov; grade 5); the House has not scheduled the Senate text and the bill's enactment odds remain the W46 20-25% (TD Cowen, Rapidan) pending that vote. Regional: six RTOs filed large-load interconnection tariffs in response to FERC's June show-cause orders; PJM's ER27-531 proposes a 12-month Fast Energization Process and a 50 MW registry (comments December 7 and January 15), and PJM's transmission owners filed Cost Recovery Agreements with ten-year minimum terms and 80-100% revenue commitments (grade 5); NIPSCO was authorised to serve Amazon's 2,400 MW by 2032 (grade 5).

State and utility: Entergy Louisiana's Waterford 5 plant for Meta's Hyperion campus is now projected at $1.755B, up $676M (63%) from $1.079B, with seven more gas plants to be filed in December (WWNO / Louisiana Illuminator, grade 3); seven PJM states' consumer advocates asked the PJM Board to seek postponement of the December 9-15 capacity auction until data-centre load is carved into the backstop, with the Independent Market Monitor projecting a clear at the $337.83/MW-day cap and $5-10B of added cost (grade 5); Loudoun County paused 33 legislative applications for 12 months and four more counties adopted moratoria (grade 5). Google signed a 50 MW superhot-geothermal PPA with Quaise for delivery from 2030 (grade 4).

Induction from the window: every instrument in it allocates cost to the load (Section 229's floor; the CRAs' ten-year 80-100% commitments; Entergy's dedicated-plant recovery; the Monitor's $5-10B at the cap) and every delivery date is 2028 or later (the BRA's 2029/30 year; NIPSCO's 2032; Quaise's 2030; Waterford 5's in-service for Hyperion), which is the pattern the re-scoped wording expects, observed again in the second window since the re-scope. The addition the week forces is the Loudoun pause, a permitting constraint arriving in 2026 in the zone that drives PJM's shortfall, which the thesis test records against the 2028 clock; Entergy's 63% overrun is recorded as the cost of generation rising, not as a measure of the constraint.

STEEL-MAN: Four bounds. First, Section 229 has passed one chamber of a Congress whose House has not scheduled it, and the W46 enactment odds (20-25%) are unchanged by a Senate vote everyone expected; the CRAs and ER27-531 are filings FERC has not accepted. Second, Entergy's 63% overrun is a single utility's revised estimate on one plant and is grade 3 reporting; gas-turbine cost inflation is an industry-wide fact that would raise the plant's cost whether or not a data centre was the customer, so it measures the cost of generation, not the cost allocated to AI specifically. Third, a capacity auction clearing at a price cap says the cap binds; it does not say how far above the cap an uncapped market would have cleared, so the Monitor's $5-10B is the cost of the cap being where it is, not a measure of scarcity. Fourth, the Loudoun pause, the four new county moratoria and the seven-state postponement request are permitting and political constraints arriving now, in 2026, which is evidence for the retired 24-month wording of the hypothesis and against the 2028 clock the re-scoping adopted; the thesis test names it as counter-evidence. The claim survives because its content is the direction and simultaneity of the allocation (three sets of proceedings, one quarter, every instrument pushing cost to the load with a 2028-or-later delivery date), which holds whether or not any single instrument is enacted, and because the house's two dated bets (p161, p167) are written against filed documents rather than legislation. Confidence is 66 for three reasons: the instances are filed or passed documents rather than reports (four of six at grade 5); the pattern has held in both windows since the re-scope (W46's Kentucky ESA and this week's six instruments) without a counter-instance that allocates cost away from the load; and the two instruments that

> would most change it (House action on Section 229; FERC's ruling on the CRAs) are pending, which caps it below 75. It happens to match p161, which tests the same pattern at the state layer.

THESIS TEST

H1 · SUPPORTED **The cycle is accelerating, not slowing: capability, density and bandwidth doubling times are shortening on every flywheel turn.** Premise: each turn should be shorter than the last on all three lenses; the falsifier is a next-generation platform below 1.5x its predecessor or any lens's doubling slope flattening for two consecutive quarters. This week's intervals: Nvidia booked about $20B of Rubin revenue in the platform's first quarter, twelve months after the GB300 ramp, while guiding the next quarter to $130B (+12% sequential) and fiscal 2028 to about 75% growth; its networking line grew 2.6x year on year and CPO scale-out reached volume production; the first Vera Rubin NVL72 racks came online at a neocloud (Nscale) with a factory-integrated networking rack. On the software lens the interval compressed in compute per token rather than index points: GPT-6.2 Astra delivers 62 index points for $4.00 per task, which at Fable 5.2's identical list price is a 6.6x step in tokens per task in nine days; Google shipped its fourth Gemini 4 model in seven weeks (Neon at 61) and took the top coding-agent slot; Inkling 2 Mini runs a 26.4B looped MoE natively in NVFP4 at 20 GB. On the hardware lens TrendForce's HBM4 range (1.5-2.7x per gigabit in 2027) and HBM at 30% of wafers are the price of a memory curve steepening; Microsoft's Maia 200 in volume production and AMD's MI430X for 1H27 are second and third accelerators on an annual cadence. The tracked metric is the interval between successive platform generations on each lens, and none lengthened. The refutation condition is not approached: Rubin's first-quarter revenue is the fastest ramp the vendor has recorded (vendor statement), and the Taiwan slip reports (grade 3) are scored as noise. Verdict history: strained W43-W44, supported W45-W47; three consecutive supported weeks. **AGAINST IT:** The TOP500 recorded its lowest turnover in list history (27 new systems, 5.4%) with the top ten unchanged, which is the HPC flywheel stalling even if the cause is hardware flowing to AI factories. Nvidia's gross-margin guide fell to 71-72% from 74.9% on the Rubin ramp, which is the cost of acceleration arriving before its volume. Astra's Coding Agent Index score did not move and its hallucination rate (59%) is a capability that did not improve with the generation. The Vera Rubin slip reports, though scored as noise, name HBM4 at 100% utilisation and a switch-board yield problem that, if either is true for any customer, would be the first platform interval to lengthen.

H2 · SUPPORTED Capital is concentrated; returns are captured one layer up the supply chain, and diffuse returns are not yet measurable (restated 2026-11-14 from the v1 'Capital is concentrated, returns are diffuse: the capex bet pays off in enterprise software, productivity and on-prem deployments, not where the capex is spent'). Second grading on the restated wording, with this week's evidence independent of the evidence the restatement was written on. Direct-capture leg (returns captured one layer up and at the labs' largest customers, with prepayment travelling up the stack): supported, at the largest scale in the series. Capital concentrated: ~$119.5B into one lab in one week, ~$44B into one lessee's GPU lease, $390B of supply commitments at one vendor, $20B raised by one hyperscaler to fund one customer's contract. Returns one layer up: Nvidia's $115.8B at a 74.9% gross margin (+103%), networking at 2.6x, Anthropic's run-rate at $81B with about 900% net revenue retention and 74% of trailing revenue from 11,900 customers above $100K. Prepayment up the stack: Saturn's $7.1B of upfront rent on ~$47.3B of racks; Nvidia's up to $105B of credit support for land, power and shell at the PORTS-Pike data-centre campus its customers' GPUs will run in. Diffuse leg (economy-wide returns not yet measurable on central-bank and statistical-agency data): supported, with the survey evidence split for the first time. For the leg: Anthropic's prospectus still has Claude under 0.1% of its Fortune 500 customers' expenses; SnapLogic has 54% of enterprises having abandoned agentic projects for unclear value and 42% able to forecast token spend; IBM has 18% of chief AI officers controlling investment; the Harvard paper on 720,542 workers has coding agents raising output measures (PRs +35.7%, lines +44%) while final deliverables (issues +13.7%, epics +18.0%) move a third as fast and review-to-merge time rises 56.1%. Against it: Bain's measurable-results share rose ten points to 33% and software teams now expect headcount growth of 10%; Intuit's agent connectors carry $1B-plus of annualised invoice volume at 2M-plus calls a week (filed). Neither leg's refutation condition (supplier margins falling while hyperscaler AI-segment margins rise; a G7 central bank or statistical agency attributing a percentage point of productivity growth to AI) is approached. The construction caveat from W46 is dropped: this week's evidence postdates the restatement. Verdict history: v1 wording strained W36-W37 and W39-W45 (supported W38), superseded 2026-11-14; restated legs supported W46-W47. **AGAINST IT:** Against the direct-capture leg: the market sold the layer-up print for a second week (Nvidia -23% on a record quarter; seventh weekly decline), the counterbrief's SemiAnalysis note (outside the graded packet) cuts the vendor's 2027-2030 shipments by more than half, and Nvidia's margin guide fell 3 points, which together say the captured returns may be peaking in the market's pricing if not in the filings. Against the diffuse leg: Bain's ten-point rise in firms seeing measurable results is the first survey in the series to move materially toward return, the software-headcount expectation swinging from -2% to +10% is the opposite of the displacement the v1 wording implied, Palo Alto paid $4.5B for a company at $100M of revenue selling post-trained open models to enterprises (a bet that diffuse returns are purchasable), and Intuit filed a dollar figure for agent-driven transaction volume. Against the grade: two of the week's four adoption datapoints came from vendors (IBM, SnapLogic) with products to sell on their findings.

H3 · STRAINED **Networking is the durable layer: interconnect holds pricing power longest because Metcalfe's law multiplies value per connected node.** Premise: the framework's test is interconnect revenue growth outpacing compute revenue growth at the operator layer for two consecutive quarters. This week the ledger corrected the first leg and the vendor layer gave its strongest print. Correction: the resolve stage scored p74 as a hit against the hypothesis on Equinix's and Csquare's primary text (Equinix interconnection $465M, +10%, against total revenue +12% normalised; the '+10% vs +7%' W44 cited compared interconnection with colocation inside stabilised assets; Csquare interconnection -7.0% against total +25.6%), so W44's 'first leg met' is withdrawn and the operator-layer test has no leg met; p154 (Q4 results by March 1) is now the first leg, not the second. Vendor layer: Nvidia's networking revenue topped $21B from $8.2B (about +156%) against compute growth of about +101% on the disclosed segments, with CPO scale-out in volume production, the first production CPO fabric disclosed at Lambda and Dell's integrated Vera Rubin networking rack at Nscale, which is interconnect content and revenue rising faster than compute at the vendor that sells both; against compute comparators in the window, Nvidia's networking +156% leads every print in the series. Pricing power is still untested rather than contradicted: Nvidia does not price networking separately inside an NVL72 rack, Eoptolink's filed denial of order cuts and 74% drawdown describe the pluggable commodity under pressure, and the FCC restriction advancing against six lobbyists treats the transceiver as a chokepointed input, not a priced asset. Standards moved toward operators specifying the fabric (IETF FANN 15-3; OpenAI's MRC; UEC's CSIG; Azure's two-tier fabric), which is the Gilder side of the hypothesis (bandwidth abundance) working against the Metcalfe side's pricing power. Verdict: strained for a seventh consecutive week (W41-W47), with the evidence relocated: the vendor leg is recorded as the strongest in the series and p166 (64%) is written to test it for a second quarter; the hypothesis moves to supported only if p166 and p154 both hold, and to refuted if p154 fails and p166 reverses. **AGAINST IT:** For the hypothesis: Nvidia's networking at 2.6x is interconnect revenue outgrowing compute at the largest vendor by 55 points, the widest margin the series has recorded on any layer; Arista launched a distributed-inference router line with coherent 400G/800G pluggables for inter-site traffic that CPO does not address; Ciena and Internet2 ran 800G coherent pluggables 2,104 km; InfiniBand's TOP500 share rose to 60.6%; and Azure's hollow-core fibre at 15,000 km with a 10x target by 2028 is an operator locking in interconnect capacity years out, which is Metcalfe's argument in physical form. The decisive test is still the operator-layer leg, and this week's correction means it has not been met once.

H4 · STRAINED Open weights pull the floor up: they are a demand catalyst for on-prem and sovereign GPU consumption, not primarily a threat to closed labs. Premise: the downloadable floor rises step-wise and the gap to the closed leader stays within ten index points; the refutation condition is a gap above ten for two consecutive quarters, now four weeks into the first quarter since the v5 ruler of October 27. Artifacts: the best week in the series. Tencent published Hy4's weights (770B / 49B active, 1M context) on November 20 and Hy4 flash's (224B / 11.3B active) on November 17-18 under the Tencent Hy License 1.0, which requires a separate licence above $50M of revenue or 100M monthly users; Thinking Machines shipped Inkling 2 Mini (26.4B / 3.4B active, Apache 2.0, native NVFP4, ~20 GB) with Unsloth and Ollama support on day one and the best sub-100B open scores on DeepSWE (48.2%) and the Vals Index (46.5%); Microsoft open-sourced Aion-1.0-Instruct-2.3B under MIT as Edge's built-in model. Score: unchanged. The best downloadable index score is still Kimi K3.1 Pro's 46 against a closed leader at 65, a gap of 19; Hy4 at 35 becomes the second-highest downloadable score and Inkling 2 Mini scores 20.1; Qwen4-Max went to production closed (2.4T / 71B active) with its open models 'coming soon' and no date. Demand route: the first large capital commitment to the route landed (Palo Alto's $4.5B for Applied Compute, which post-trains open weights for enterprises; Model ML's claim of 2% of frontier cost on open models), and the Hugging Face harness-adapter result (a 320 MB LoRA lifting Qwen3.8-27B by 14 points on SWE-bench Pro and Terminal-Bench 2.1 at 3.5-4.4x fewer tokens; Agent Techniques Weekly) is the technical case that open models plus a harness close part of the gap in production even if not on the index. Against the route: the BIS remote-access rule licenses sovereign capacity in Southeast Asia serving Chinese customers, and OpenAI's anti-distillation controls (54.0% to 0.2%) are the model-layer half of a regime that raises the cost of fast-following for the labs that produce most open weights. Verdict: strained (W37-W38, W40-W47); moves to refuted if the gap is still above ten at the end of the second quarter in late April, to supported if a downloadable model lands within ten points of the leader before late January; p164 (28%) is the dated bet on 50 by January 31. **AGAINST IT:** For the hypothesis: the demand route got its first $4.5B cheque and its first production-grade technique (harness adapters) in the same week, Hy4's weights landed at the previewed architecture after sixteen days rather than never, and the local hardware floor moved to 20 GB for a model that leads its size class. Against it beyond the gap: the Tencent licence's revenue and user thresholds mean 'open weights' and 'deployable by a large enterprise' have come apart on a licence term rather than a territory for the first time; the Hy4 system card's monitorability result (40.59% hidden / 0.00% visible) is a property enterprises will have to manage in the model they can download; and the counterbrief's Interconnects reading (credited) that the US open ecosystem has lost its political cover is consistent with the anti-distillation and remote-access rules landing in one week.

H5 · SUPPORTED **Power and permitting are the binding constraint on capacity from 2028; through 2027 the constraint is advanced packaging, HBM and substrates, and the cost of power is being allocated to large loads by regulators before it binds physically (re-scoped 2026-11-14 from the v1 'Power is the binding constraint for the next 24 months: no longer chips, capital or land, but electricity, visible as PPAs, behind-the-meter generation and interconnection queues').** Second grading on the re-scoped wording, on evidence that postdates it. Components-through-2027 leg: supported in the buyer's filing and the forecaster's range. Nvidia raised supply commitments to $390B from $279B with multi-year memory procurement the largest component and $203B due in fiscal 2028, and signed substrate and packaging deals with Kinsus (3.5x capacity), SEMCO and SPIL; TrendForce priced 2027 HBM4 at $3.00-5.50 per gigabit (1.5-2.7x) with HBM taking 30% of DRAM wafers; the Vera Rubin slip reports, scored as noise, name HBM4 supply and switch-board yields rather than power as the constraint they allege. Power-from-2028 leg: supported in every delivery date in the window: PJM's December auction is for the 2029/30 year, NIPSCO's Amazon service is 2,400 MW by 2032, Google's Quaise PPA delivers from 2030, Entergy's Waterford 5 serves a campus whose in-service is 2028 or later, and PJM's Fast Energization Process is a 12-month path for loads that bring their own capacity. Cost-allocated-before-it-binds leg: the week's strongest, with instances at three layers (the fourth synthesis connection): the Senate passed Section 229 with no megawatt floor 62-33; PJM's transmission owners filed ten-year 80-100% revenue-commitment CRAs; Entergy's dedicated plant for one tenant rose 63% to $1.755B with seven more to be filed; the Independent Market Monitor projects the capacity auction at its cap with $5-10B of added cost that seven states' advocates want carved out. The refutation conditions (two consecutive quarters of supplier prints naming power rather than packaging or memory as the 2026-27 constraint; a grid expansion shortening queues within twelve months; a 10x inference-efficiency breakthrough) are not approached; Cerebras's vendor-claimed 5x tokens per megawatt and Maia 200's 1.7x tokens per watt are efficiency claims without baselines and are not the 10x. The W46 construction caveat is dropped. Verdict history: v1 wording strained W33, W36, W37, W40-W45, superseded 2026-11-14; re-scoped legs supported W46-W47. **AGAINST IT:** Against the 2028 clock (and for the retired 24-month wording): Loudoun County, the densest data-centre market in the world, paused Board action on 33 legislative applications for up to 12 months on November 17, and four more counties adopted moratoria in the same week, which is permitting binding in 2026 in the zone that drives PJM's shortfall; the seven-state request to postpone a capacity auction two weeks out is a political constraint arriving on the 2026 calendar; and Nvidia's up to $105B of credit support for land, power and shell at SB Energy's 4.25 GW PORTS-Pike data-centre campus is the component vendor guaranteeing the buildings and the power behind its customers' racks, which is the vendor saying the site, not only the chip, is what it has to backstop now. Against the components leg: Nvidia booked about $20B of Rubin revenue and called it the fastest ramp ever, which is a component constraint not binding the one buyer that committed its way through it, and the margin guide, not the volume, is where the constraint shows. Against the grade: the Entergy figure is grade 3, and the Senate vote was the expected outcome of a cloture the W46 issue already recorded.

PATTERN WATCH

Agent-model list prices fall or converge while the effective bill moves through non-list meters: effort tiers and tokens per task, cache rates, context breakpoints, speed tiers, dated promotions with published reversions, credit pools, subscription-only access, per-agent subscriptions, and per-action and per-resolution tolls. (12 WEEKS OBSERVED)

Next week: Confirming next week: Google does not extend Carbon's or Argon's introductory rate before Dec 1 and the $4 / $20 step takes effect on Google's pricing page; Alibaba retires the Qwen4-Max preview on Nov 25 without a list cut on the production API; and no lab in the $2 / $10 band cuts list. Breaking: Google cancels or extends the Dec 1 step; a lab in the band cuts list outright; OpenAI moves dots' Jan 14 billing date; or any major provider publishes a vendor-authored worked bill for a full agent session itemising effort, cache, context, tools, per-agent subscription and resolution charges on one page.

The label and the artifact keep coming apart, now narrowed: after the Western-lab leg closed for a fourth week and the Chinese-lab leg closed this week (Hy4 and Hy4 flash landed as previewed), the pattern is restated as vendor measurements published without a named baseline or an independent run, plus Alibaba's Qwen4 as the one open-weights promise outstanding. (12 WEEKS OBSERVED)

Next week: Confirming next week (narrowed pattern): no Qwen4 open-weight repository appears by the Nov 25 preview end; no vendor publishes a named-baseline, power-measured figure behind the Maia 200, Cerebras or Intel claims; and no InferenceX Rubin result appears before Dec 1. Breaking: a Qwen4 repository lands with weights and a licence; Microsoft names the GPU and workload behind 1.7x tokens per watt or an independent lab reproduces it; or Nvidia or a Rubin customer (Nscale, SpaceX, Lambda) publishes an InferenceX result with cost per million tokens. If the narrowed pattern holds for two more weeks, it should be renamed 'vendor efficiency claims without baselines' and the label-versus-artifact series closed as resolved.

AI returns print one layer up the supply chain from where the capex is spent (foundry, memory, equipment, packaging, ODM assembly, optics, switching) and at the labs' own revenue lines, while the enterprise-software layer the framework expects to capture diffuse returns is punished on consumption metrics and macro measurements of diffuse productivity stay near zero; since W46 the equity market has stopped paying for the layer-up prints. (7 WEEKS OBSERVED)

Next week: Confirming next week: any supplier printing in the Thanksgiving week (Dell's Q3 FY27 is the scheduled one) reports AI-infrastructure revenue growth above 50% year on year and its stock closes the week below its pre-print level; CREFC records a fifth week without a data-centre securitisation; and no hyperscaler or lab discloses an AI-segment operating margin. Breaking: a supplier guides 2027 AI revenue or margins down; an enterprise-software vendor is rewarded by the market for a consumption print; ANT closes above its $111.50 opening print; or the BLS or a G7 central bank revises measured AI productivity materially upward.

SECOND-ORDER EFFECTS

Nvidia's balance sheet becomes the allocator of 2027 memory and a guarantor of the campuses its customers' racks will run in, which changes what every other accelerator buyer pays and where. If one buyer has committed for the largest share of 2027 HBM4 output, the marginal bit for AMD's MI430X and MI455X ramps, Google's TPU programme (the reported direct-HBM purchases and the Kuo report of a TPU 8i cut, both low-grade, are consistent with a buyer looking for a way around the order book), Microsoft's Maia and Amazon's Trainium clears at the top of TrendForce's range rather than the bottom, and the custom-silicon cost advantage over Nvidia narrows by the HBM premium those buyers pay; the house's second-order read of W46 (that Nvidia's HBM discount was personal to its CEO) is now moot because the discount, if it exists, is contractual at $390B. On the site side, a component vendor providing up to $105B of credit support for the land, power and shell of a 4.25 GW campus leased for 20 years to one lab is the vendor-backstop W46 said had been syndicated reappearing one layer down the stack: not a GPU floor for a neocloud but a credit wrap for the buildings and the power behind them, which puts Nvidia's rating behind a data-centre developer's financing and makes its own cost of capital a function of a multi-gigawatt build and of OpenAI's lease performance. Trigger: Nvidia raised its supply commitments to $390B from $279B in one quarter, with multi-year memory procurement the largest component and $203B due in fiscal 2028, signed substrate and packaging deals with Kinsus (3.5x capacity), SEMCO and SPIL, and disclosed $118B of guarantees including up to $105B of credit support for land, power and shell at SB Energy's PORTS-Pike campus, a 4.25 GW, nine-building data-centre campus in Ohio under 20-year OpenAI leases with the first building due in fiscal 2029; TrendForce the next day priced 2027 HBM4 at $3.00-5.50 per gigabit (1.5-2.7x) with HBM taking 30% of DRAM wafers. Horizon: Micron's FQ1 2027 print in December and SK hynix's and Samsung's Q4 calls in late January for the first HBM4 contract disclosures (p165); Nvidia's Q4 FY27 in late February for the next commitments figure; SB Energy's PORTS-Pike financing close, date unannounced, with the first building due in fiscal 2029.. Who moves: AMD, Google, Microsoft, Amazon and Broadcom as competing HBM buyers; SK hynix, Samsung and Micron pricing desks; Kinsus, SEMCO and SPIL as newly committed suppliers; SB Energy, OpenAI as the campus tenant, and the utilities interconnecting PORTS-Pike; the rating agencies that now have to treat Nvidia's guarantees as contingent liabilities; and the neoclouds whose 2027 capex per megawatt moves with the HBM4 clearing price..

The constraint on Chinese labs moves from chips to compute access to capability transfer, and each layer lands on a different party. The remote-access rule lands on Southeast Asian and Brazilian data-centre operators, including the sovereign vehicles the on-prem row has charted (W46's Sarawak vehicle; Gulf operators serving Asian customers), who now need a US licence to serve Chinese tenants above a compute threshold and have ninety days to get one; the ECRA disclosure lands on Nvidia and AMD, whose China licences are now a public number Congress can cap in the NDAA conference; and the anti-distillation layer lands on the open-weight ecosystem the fourth hypothesis depends on, because the labs that ship most open weights are the labs whose fast-following cost just rose, and the controls that stop extraction from a closed API also make a closed API harder to use as a teacher for an open student. The house's read is that the three layers together make 'where is the model' a licensing question at every level, and that the open-weights gap (19 points) is now partly a policy variable: a Chinese lab that cannot distil from the closed frontier closes the gap more slowly, and a US lab that can is the one a buyer should expect to move the open floor next. Trigger: In one week the export-control regime acquired a third layer: BIS's interim final rule on remote access (effective November 19; licences for advanced-chip data centres in Singapore, Malaysia, Thailand, Indonesia and Brazil serving Chinese customers above 12M aggregate TPP; the Diffusion Rule rescinded; comments January 19; a 90-day licence window to February 17) controls where compute can be reached from; the House Foreign Affairs Committee's ECRA disclosure that BIS approved 879,868 H200 and MI325X chips for Chinese firms out of 3,444,160 applied for controls what ships; and OpenAI's anti-distillation controls (adaptive extraction of protected reasoning cut from 54.0% to 0.2%; reasoning bound to conversation history; thousands to tens of thousands of accounts enforced per week) control what can be learned from the model's outputs, with the counterbrief's HSBC note on Zhipu (a two-thirds price-target cut citing US labs' anti-distillation measures; outside the graded packet) as the first market read. Horizon: December 31 NDAA conference (Chip Security, AI OVERWATCH, MATCH riders); January 10, 2027 US and China suspension cliffs; January 19 BIS comment deadline; February 17 end of the 90-day licence window; MOFCOM's promised countermeasures, undated.. Who moves: Data-centre operators in Singapore, Malaysia, Thailand, Indonesia and Brazil and the sovereign funds behind them; Nvidia and AMD's China licence books; Chinese frontier labs (Zhipu, Alibaba, Tencent, Moonshot, DeepSeek) and the enterprises that deploy their open weights; US and European labs releasing open weights; hyperscalers with Asia-Pacific regions serving Chinese customers; and the NDAA conferees..

Marginal AI demand gets pushed toward lease-form compute and away from owned, grid-connected campuses, because the two markets now price the same risk in opposite directions. A campus owner carries a ten-year power commitment at 80-100% of projected revenue plus a share of a dedicated plant whose cost just rose 63%, financed by construction debt that trades in the mid-80s if the anchor tenant serves notice; a lessee of racks under a hell-or-high-water lease carries a five-year obligation financed at 8.259% by investors who underwrite only the lessee's signature, with the power obligation sitting with whoever hosts the racks. In Saturn's case the lessee and the host are the same party (SpaceX leases the racks and runs them), so the split the house describes is a hypothesis about the next lessee, not an observation about this one. The arbitrage, if it holds, is that the lessor's investors have priced the racks' depreciation and not the site's power commitment, so the power cost the regulators are allocating lands on the host, who in the neocloud model is a different party from the lab or the lessee. The house's read is that the regulatory allocation and the credit sort are two halves of one shift: cost is moving to the load, and the load is moving to a lease. Trigger: The fixed-cost floor for a US gigawatt campus rose and lengthened in one quarter: the Senate passed Section 229 (loads bear their own costs; no megawatt floor), PJM's transmission owners filed ten-year Cost Recovery Agreements with 80-100% revenue commitments, Entergy's dedicated plant for one campus rose 63% to $1.755B, and the capacity auction is projected to clear at its cap with $5-10B of added cost, while the same week's credit market cleared GPU-lease paper (Saturn at 8.259%) and agreed hyperscaler-guaranteed shells (the AWS portfolio at a 7.0% cap) and saw a construction loan slice trade in the mid-80s (Jupiter, with an investment-grade tenant under a force-majeure notice). Horizon: December 7 and 16 FERC comment deadlines on ER27-531 and the show-cause responses; December 9-15 PJM BRA (p167); Entergy's December filings for seven more plants; January 15 ER27-531 second comment round; the House's action on H.R. 2347, unscheduled; p161 by March 31.. Who moves: Neoclouds and developers who host leased racks and carry the power commitments (CoreWeave, Nscale, Lambda, Crusoe); utilities writing ten-year CRAs and dedicated-plant recoveries (PJM's transmission owners, Entergy, NIPSCO); state commissions that inherit Section 229's standard if enacted; GPU lessors and their rating agencies (Apollo, KBRA); and hyperscalers deciding whether to own the next gigawatt or lease the racks that fill it..

STRATEGIC OUTLOOK

Read the week by signature. If you run capital: the AI trade has a credit curve, and this week it ran from 5.95% (the IG index yield) through 7.0% (the cap rate on Amazon-guaranteed shells), 8.259% (the yield on a SpaceX rack lease, A-, a spread continuous with October's EdgeCore print) and 7.00-8.65% (Oracle unsecured, W45) to the mid-80s cents on the dollar (a slice of a construction loan on a lab-anchored campus with an investment-grade tenant under a force-majeure notice) and two conduit deals quoted for four weeks without a print, plus a neocloud loan with no pricing reported through November 22. Place every commitment you hold on that curve by the form of the contract behind it and how enforceable it is during the build, not the asset, because that is how the market sorted it (62% confidence; the unsigned half rests on two grade-3 facts and an absence of reports). Use the house measurement as the template for stabilised sale-leasebacks: at a 7.0% cap on 12-year paper, 67% of the gross value is contracted rent at 5.95% and 33% is a 2038 residual, carried for about 105 bp over the index, with leverage adding 0-100 bp on a redacted coupon. ANT is the public mark on the labs' commitment chain; it touched $83.05 and closed day two at $88.91, 7% over issue, on a day the Nasdaq fell 5.3%. Stress to a close below $83 (p168, 57%) and to a December rate hike, not to the IPO's reception. If you run infrastructure: the public mark for hyperscaler shell is now $7.7M per IT-megawatt and $45 per kW-month on a 12-year Amazon guarantee, and anything turnkey must clear well above it; the 2027 memory you need is in Nvidia's $390B order book, so plan allocation before price and expect HBM4 at the top of TrendForce's $3.00-5.50 range if you are not the largest buyer (p165, 61%); your 2028 power cost is being set this quarter in four proceedings that all allocate it to you (Section 229, which still needs the House; PJM's CRAs and Fast Energization at FERC; Entergy-style dedicated plants at state commissions; a capacity auction at its cap), and the Loudoun pause says permitting can bind in 2026 even where the thesis dates power at 2028. Sign for energised megawatts; if you host leased racks, the power commitment is yours and the lessor's investors have not priced it. If you run an AI platform: the frontier token got 6.6x cheaper per task at the same list price and maximum effort (2.3x against Fable's high setting) by moving reasoning out of the visible stream, and the outside monitor that could not read it is the one you are probably running (58%); contract per completed task at a named effort level and put the current-generation monitor or the reasoning-visible mode in the same budget line. The open floor had its best week for artifacts (Hy4, Hy4 flash, Inkling 2 Mini, Aion) and did not move in score (46 against 65); keep the routing path, read the licence thresholds, and treat the harness-adapter result (a 320 MB LoRA lifting a 27B open model 14 points at fewer tokens) as the production-grade way the gap closes before the index does. If you run a network: the fabric is being sold inside the rack and financed inside the lease, Nvidia's networking outgrew its compute (p166, 64%), and the operator layer the third hypothesis is written on has not met its first leg once; dual-source 2027-28 pluggables against an FCC restriction six hyperscalers and vendors could not stop, and treat CPO as a switch-vendor feature that moves the optics bill of materials. If you sit on a board: two hypotheses restated last week graded supported again on evidence that postdates them (H2's direct-capture leg at the largest scale in the series; H5 with cost allocated to the load in three sets of proceedings in one quarter), one stayed strained with its evidence relocated from the operator layer to the vendor layer (H3), one stayed strained on a 19-point gap in the best artifact week yet (H4), and the cycle is still accelerating (H1). The returns print one layer up and the market sold them for a second week; the money cleared where one investment-grade name signed for the rent. Hold the house's connections at the confidence they carry and expect the dated tests (p157 and p168 by December 31; p164 and p162 by January 31; p163, p165 and p166 by February 28; p160, p161 by March 31) to do the grading.

WHERE WE DIFFER

Our read against the field's published takes on the same week — on the record, so you can score us later.

EXTEND Market Briefer, 'The Global Close, November 20, 2026': Friday was a valuation unwind, not a credit event: the 10-year fell to 4.91% while tech dropped 7%, high-yield at 347-352 bp describes repricing rather than a funding problem, and the same day's $32B SpaceX print shows capital is available at a price that charges properly for risk.

Our read: Correct for liquid high-yield and for anything with an investment-grade signature behind the rent, which is what cleared or was agreed (Saturn at 8.259%, with ~$26B pre-committed; the AWS portfolio at a 7.0% cap, agreed November 11). The house adds the half the frame leaves out: a slice of a construction loan on a lab-anchored 2.45 GW campus with an investment-grade tenant under a force-majeure notice traded in the mid-80s, no data-centre ABS or CMBS has priced in four weeks while two deals sit quoted, and a neocloud's investment-grade-structured loan had no pricing reported through November 22. For that paper the week was a credit mark on the enforceability of the signature during the build, and the market has sorted by signature rather than repriced as a whole; p163 (38%) is the bet that the sort lasts a quarter.

DIFFER Matt Levine, Money Stuff (Bloomberg Law), November 18-19; Zvi Mowshowitz, AI #195 (November 19); Morningstar (November 19): Anthropic's IPO was underpriced: Levine set up the $100-versus-$120 pricing question on November 3, noted the book six times covered at the original range and that bankers 'remain remarkably close to undefeated at convincing companies to underprice' (November 18-19); Zvi Mowshowitz wrote that $83 'still seems obviously too low' (AI #195, November 19); Morningstar put fair value at $92.

Our read: The pop reverted inside two sessions: day two closed at $88.91, 7% above issue, with an intraday low of $83.05, on a day the Nasdaq fell 5.28% and the semiconductor index 11%; the counterbrief's Blockworks item (outside the graded packet) has an offshore perpetual-futures market near $95 all week, which makes the $111.50 open the anomaly and $83 a price the market converged toward, not away from. The house reads the stock as the public mark on the lab's roughly $895B of commitments; a close 7% over what seven suppliers paid, with the low a nickel above it, says those commitments are priced near the suppliers' terms, and the tape that day is a confounder the house cannot remove. p168 (57%) bets it trades through issue by year-end.

OPEN CNBC, Nvidia earnings live blog (November 17) and seven-week losing streak (November 20); Investing.com and Motley Fool for the 'gone too far' framing: Nvidia trades at roughly 8.4x forward earnings after the November 20 drop against AMD near 29-30x, 'very cheap relative to its peers' on FactSet's numbers (CNBC, November 17 and 20); the 'absurdly cheap' and 'selloff has gone too far' readings are the retail-commentary layer (Investing.com, Motley Fool, November 20-21) built on CNBC's multiple.

Our read: The multiple is cheap only if the denominator holds, and the counterbrief's strongest hardware argument (SemiAnalysis's November 13 model cutting 2027-2030 shipments by 22/39/56/66% and Feynman pricing from ~$110K to ~$78K, explicitly independent of the CEO change; outside the graded packet) is a dispute about that denominator the print did not settle. The house's own numbers cut both ways: $390B of supply commitments is a buyer committing to a 75% growth year, and a 71-72% margin guide is the Rubin ramp's cost arriving first. The house holds no view on the stock and records the question as open until the Q4 print and the first supplier HBM4 price (p165, p166).

DIFFER Simon Willison, 'GPT-6.2 Astra and Sol' (November 17); Artificial Analysis release note (November 17); Zvi Mowshowitz, AI #195: GPT-6.2 Astra priced at exact parity with Claude Fable 5.2 ($8 / $40), with the 60% cut to cached input 'great news for longer running agentic sessions' and 'the buzz around these new models is mixed' (Simon Willison, November 17); Artificial Analysis has it second on the index (62 against 65) with the Coding Agent Index flat against GPT-6 Astra (November 17); Zvi called early user feedback 'unpromising' (November 19). The counterbrief's composite of these is 'a price cut, not a ceiling'.

Our read: At identical list prices the per-task bill is $4.00 against $26.54 with both at maximum effort, a 6.6x gap that is almost entirely tokens per task (2.3x against Fable's high setting), and the system card itself attributes lower monitorability at comparable effort to that token efficiency, a reading the counterbrief's Epoch and Ground Truth items share. That is a ceiling moving in compute per token rather than in index points, and it came with a measured cost: 48.8% evasion of a prior-generation action-only monitor against 0.6% for a current one. Willison's cache point is the same fact from the buyer's side: the token is cheaper to re-read and harder to read. The house's read is that the buyer who takes the saving has to buy a current-generation monitor with it; The Model Pulse carries the numbers.

EXTEND Nathan Lambert, Interconnects, 'Open-source AI without Jensen' (November 14); Zvi Mowshowitz, AI #195 (November 19); Talking Fingers (November 21): Lambert (November 14, before Inkling 2 Mini and the Hy4 weights) asked whether Western open models can keep up without Nvidia's backing ('maybe they can pull it off?'); Zvi (November 19) then wrote that there is 'a very large gap between open and closed models, and we still have not seen a Fable class open model', that Qwen4-Max went closed and sits 'well below the cost-capability frontier', and that Inkling 2 Mini is the best Western open model in a while but far behind Kimi K3.1 Pro; Talking Fingers (November 21) carried the Hy licence's $50M / 100M thresholds.

Our read: Agreed on the score (the best downloadable model is still 46 against a closed leader at 65) and the house adds the two things the frame understates: this was the best week for open artifacts in the series (Hy4 at 770B and Hy4 flash as previewed, Inkling 2 Mini as labelled with day-one Unsloth and Ollama support, Microsoft's Aion under MIT), and the first $4.5B acquisition whose thesis is selling post-trained open weights to enterprises (Palo Alto and Applied Compute) landed the same week. The fourth hypothesis stays strained on the gap; p164 (28%) is the bet that the gap closes to 15 by January 31, and the licence-threshold clause is the new line on the open-weights checklist.

EARLY WARNING PANEL

The levers we monitor.

14 metrics, current vs prior period. **2 rising**, **0 falling**, 12 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 (Opus); the dated crossing continues: Gemini 4 Carbon at $0.10 cached through November 30, reverting to $0.20 on December 1; GPT-6.2 Sol holds $0.10 cached at list; Neon enters at $0.15; Qwen4-Max at $0.075**	$0.20 (Opus); threshold crossed by another model for a dated window: Gemini 4 Carbon in production on the paid API at $0.10 cached (implicit caching, 95% off $2) from November 10 through November 30, reverting to $0.20 cached on December 1	→	Cache reads at or below $0.10 on a frontier model in production

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Mid-tier frontier list price, labs at $2 input / $10 output	**Five models from four labs at list: Claude Sonnet 5.5, GPT-6.2 Sol (replacing GPT-6.1 Sol at the same price), Gemini 4 Argon and Gemini 4 Carbon (both stepping to $4 / $20 on December 1), Kimi K3.1 Pro. Between the band and its double: Gemini 4 Neon at $3 / $15. Below the band: Qwen4-Max at $1.50 / $4.50 (closed), Gemini 4 Hydrogen at $1.50 / $7.50, GPT-6.1 Terra at $1 / $5**	Five models from four labs at list: Claude Sonnet 5.5, GPT-6.1 Sol, Gemini 4 Argon and Gemini 4 Carbon (both introductory, stepping to $4 / $20 on December 1), Kimi K3.1 Pro (downloadable weights). Below the band: Gemini 4 Hydrogen at $1.50 / $7.50, GPT-6.1 Terra 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	**8.259% (Saturn Funding Trust A3, $25.168B, KBRA preliminary A-, due April 2033, T+345), the new high; Lambda's 6.78% (A(low) / Baa1) is now the low. Guided and with no pricing reported through November 22: CoreWeave's $4B-plus Meta-backed IG-structured loan at 7-7.5% talk, first reported November 10 (Semafor, November 13)**	6.78% fixed (Lambda $1.008B, A(low) / Baa1), no new rated GPU-backed print; the next one is guided: CoreWeave expects a 'high investment-grade' asset-level facility within one to two weeks at ~7-7.5% (CEO on CNBC, grade 5 for the guidance, no term sheet). Unrated comparison prints this week: WhiteFiber $300M at SOFR+600 with a 2% floor (about 10% at current SOFR) for an Nscale-leased campus; Blue Owl's $512.6M floating-rate mortgage on a hyperscaler-leased building, rate undisclosed	↑	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%), no new Volta quote; the wider distribution moved out: US high-yield OAS 343 to 352 bp, CCC ~1,316 bp, IG ~90-94 bp; the Jupiter construction loan traded in the mid-80s (grade 3); no new CoreWeave 2031 quote**	~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 distribution narrowed at the top for the first time in four weeks: CoreWeave's 2031 notes fell to a 14.79% yield from 16.7% after its November 10 print (CNBC, November 13; outside the graded packet), ~800 bp over Lambda from ~990 for the same 2031s a week earlier (W45 quoted the 2032s above 16%; the series now stays on the 2031s, the note with a print); Galaxy's Helios data-centre bonds at 89.25 to yield ~13% against a 9.875% July coupon (B+); WhiteFiber's	→	Spread below 150 bp (collateral is being underwritten) or above 500 bp (counterparty risk is being repriced)

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
		new unrated loan at SOFR+600 (~320 bp over Lambda); US high-yield spread 348 bp on November 5 easing to 342 bp, single-B 367 bp		
Broadcom five-year CDS, bp	**No absolute print in the graded packet for a third week; carried as 'up, level unconfirmed' from W45's +49% monthly move (house arithmetic on the W42 base gives ~200 bp, the threshold). This week's Broadcom facts: shares -11-12% on the Kuo TPU 8i report (grade 2); Class N equity in Anthropic at $83 now marked at $88.91**	No absolute print in the graded packet for a second week; carried as 'up, level unconfirmed' from W45's CNBC-reported +49% monthly move (house arithmetic on the W42 base of 133-136 bp gives ~200 bp, the threshold). This week's Broadcom facts cut toward the exposure, not away from it: Broadcom is one of the seven Class N buyers in Anthropic's $26B placement (amount per buyer undisclosed) and its shares rose ~6% on November 12	→	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	**Superseded in practice, not lifted in words: GPT-6.2 Astra shipped to the API, ChatGPT Work and Codex on November 17 with tool use, rated Critical for cybersecurity and High for AI Self-Improvement (the first external deployment at that level), 62 days after the September pause; OpenAI's November 18 'Managing frontier AI risk' post names GPT-6.1 Astra as a model it held back and describes pausing and holding back as ongoing practice. No first-party statement that the pause is lifted**	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 (55 days). No first-party resume statement; the week's OpenAI shipments were harness and distribution: Codex CLI 0.168 / 0.169, ChatGPT Enterprise in AWS Marketplace, Businesses in ChatGPT, and a fix after Depi's disclosed prompt-injection chain through openai/codex CI	↑	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 formally revised; the roadshow adds a per-gigabit range for HBM4 of $3.00-5.50 in 2027 against $2.05 in 2026 (+46% to +168%) and HBM4E above $5, with HBM taking ~30% of DRAM wafers (from 23%), HBM bits +65.7% and PC DRAM bits -7.9%. Conventional DRAM contract prices +10-15% QoQ in 4Q26, then +5-10%, +3-8%, 0-5%, 0-5% through 2027**	+121% (not revised). Demand-destruction leg confirmed on the consumer side: notebook shipments -15.8% QoQ in 3Q26, PC DRAM contract prices +13-18% in 4Q26 (decelerating as OEMs cut procurement), enterprise SSD +15-20%, some server memory booked through 2028; Korean suppliers reportedly planning >20% cuts to 2027 PC DRAM output to feed servers (TrendForce October 30 contract report, via the counterbrief; outside the graded packet). HBM4 price reports: SK hynix mid-$500s per 12-high stack, Samsung seeking parity, yield above 70% (Korean press, single-source, grade 4); no	→	A revision below +80% or a contracted per-Gb price disclosed by a supplier

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
		per-gigabit contract price disclosed by a supplier		

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
PJM Reliability Backstop Procurement (RBP) effective date	**ER27-454 unchanged: comments due November 30; order requested by January 29, 2027; procurement to open March 1, 2027. New alongside it: PJM filed ER27-531 (Fast Energization Process, 12-month path for loads that bring or curtail capacity; 50 MW Large Load registry) on November 16 with comments due December 7 and January 15; PJM transmission owners filed ten-year minimum Cost Recovery Agreements with 80-100% revenue commitments; six RTOs answered FERC's June show-cause orders (comments December 16). Federal: the Senate passed**	Re-docketed November 9 as a Federal Power Act section 205 rate filing in ER27-454-000, 'substantively identical' to ER27-330 (to be withdrawn); comments due November 30; order requested by January 29, 2027; most revisions effective February 28, 2027 with Electric Distributor large-load reporting (Attachment DD §18.3A) effective January 29; procurement to open March 1, 2027 with a four-week sell-offer window; illustrative RBP Credit Rate (the collateral load-serving entities post against their backstop obligation) cut to ~$686,000/MW (from $1.5M/MW) at the $555/MW-	→	A PJM refile accepted without suspension, or the suspension extended past February 2027

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
	H.R. 2347 62-33 on November 19 with FPA Section 229 (no MW threshold; prospective) and a transmission-pricing NOPR within 90 days; the House has not scheduled it	day UCAP cap; 15-year commitments from the 2028/29 delivery year. Federal context: SA 6883 deletes the 20 MW covered-load floor and adds FPA Section 229; cloture 5:30 p.m. November 16 at 20-25% enactment odds		

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
PJM data-center load cancelled or delayed since the capacity-auction inputs were set	**4,017 MW (42% of forecast data-centre growth), reiterated by the Independent Market Monitor in the seven-state advocates' November 18 letter; no new figure. Outside PJM the moratorium map reached the largest market: Loudoun County paused Board action on 33 legislative data-centre and substation applications for up to 12 months (retroactive to September 15); new county moratoria in Pitt NC (two years), Washington AR (12 months), Lyon KS and Saguache CO; Broward FL votes December 8**	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 stands in the ER27-454 record. Outside PJM, the moratorium map grew by four counties and one took effect: Lee County GA to June 30, 2027; Meade County KY to January 13, 2028; Broward County FL first hearing on a 50 MW moratorium (vote December 8); Missoula County MT interim zoning study approved; St. Johns County FL effective November 15. ERCOT's Batch Zero reports remain due December 10; no new governor-elect directives in Texas, Ohio or Georgia	→	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	**None. Nvidia's Q3 call cited about $20B of Rubin revenue in the platform's first quarter and the 'fastest ramp ever' with no InferenceX result and no cost per million tokens; the W43 'official preview' (1.81-2.03x tokens per TCO dollar, modeled, no valid power telemetry) stands. The competing claims this week are also unmeasured: Maia 200 at 1.7x tokens per watt against an unnamed GPU baseline, Cerebras disaggregated inference at 5x tokens per megawatt against its own CS-3-only configuration**	No new InferenceX result; the W43 'official preview' (1.81-2.03x tokens per TCO dollar at 100 tokens/s per user, modeled TCO, no valid power telemetry) stands. The competing claim this week is AMD's: a vendor-measured 21,333 tokens/s per GPU on DeepSeek-R1 for MI455X said to be 30% above Nvidia's audited MLPerf v6.1 Vera Rubin entry (6.1-0106), no batch size, latency target or power disclosed. Rubin volume timing: Quanta ships small volumes this quarter with meaningful volume in Q1 2027; Nvidia reports Q3 FY27 on November 17 under an interim CEO	→	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). No handset event this week; the on-device datapoint is a browser one: Microsoft's Aion-1.0-Instruct-2.3B (MIT) becomes Edge's built-in model, replacing Phi-4-mini, with a Prompt API origin trial; the desktop datapoint is Inkling 2 Mini at ~20 GB in native NVFP4 (26.4B / 3.4B active)**	30B-parameter mixture-of-experts on Snapdragon 8 Elite Extreme Gen 6, vendor claim (no independent result). No handset event this week; the local-inference datapoints were desktop and vendor- or author-measured: Edge0-145B (10.12B active, Qwen3.8-Flash-Next derivative) at 98.3 tok/s on one RTX 5090 D with 23 GiB of VRAM plus 38 GiB of host RAM; Boole's 1.8-bit Flint at 565.8 tok/s on one RTX 5090 via a llama.cpp fork; llama.cpp v1.0.0 with FP8 KV caches and a host-memory MoE expert cache	→	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). Demand side: Nvidia networking above $21B with CPO scale-out in volume production and Lambda's 10,000-GPU production CPO fabric disclosed; Eoptolink's filed denial of order cuts with 1.6T in volume ramp and NPO shipments in 2027; Arista's Etherlink R4 routers with 400G/800G coherent pluggables. Regulatory side: Reuters reports the FCC transceiver restriction advancing despite Amazon, Google, Meta, CoreWeave, Cisco and Microsoft lobbying; the Covered List as of November 18 names no**	Early 2029 (unchanged); the acquirer's equity halved while the horizon held: Marvell's CEO told employees in a Form 425 filing that Marvell shares fell about 51% over the month alongside a 49% average decline for AI-connectivity peers, with Lumentum down similarly, and reiterated the all-stock combination closing in H1 calendar 2028 with an S-4 to follow. Demand side: Cisco won a 1.6T client-side optics design in the quarter and has hyperscalers running Acacia coherent optics in production; Coherent signed a letter of intent for at least $601M of 6-inch InP substrates for 2027-29 from a Chinese supplier,	→	The horizon moving in by two quarters, or a second optics vendor reporting a cancellation

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
	transceiver maker	conditional on an export licence; no optics vendor reported 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.

6 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

38%

BY FEBRUARY 28, 2027

By February 28, 2027 at least one US data-centre ABS or CMBS transaction of $300 million or more prices whose collateral is not majority-leased to, or guaranteed by, a hyperscaler or other investment-grade-rated tenant (that is, a multi-tenant, neocloud-leased, construction or development-stage pool), and CREFC or the rating agency's presale reports it as priced, that is, the conduit market reopens for the unsigned half of the bifurcation the big story describes within a quarter.

TRIGGER: Hit if, by the deadline, CREFC's CMBS / ACMBS and Data Center New Issue Weekly, a rating-agency presale or final report (KBRA, Moody's, S&P, Fitch, DBRS), or an SEC filing records a priced US data-centre ABS or CMBS of $300M or more whose disclosed collateral is less than 50% leased to or guaranteed by hyperscalers or investment-grade-rated tenants (the resolver uses the presale's tenant table; a neocloud tenant counts as non-IG unless it holds a public IG rating); a priced deal of $300M or more that is 50-75% IG-tenanted, or a non-IG-tenanted deal below $300M, is a partial; no such pricing by the deadline, or only hyperscaler-leased deals (Blue Owl's Amazon SASB, any AWS or Microsoft-leased pool) pricing, is a miss. The bet sits below even because CREFC has recorded four straight weeks without any data-centre ABS or CMBS pricing while two deals (Compass Datacenters' $380M CMPDC 2026-1 ABS and the $309M MSC 2026-ROYL Kansas City CMBS) have sat in the market unpriced, because the one observed trade in construction risk this week was HSBC's Jupiter sale in the mid-80s, because CoreWeave's investment-grade-structured loan had no pricing reported through November 22 (Semafor, November 13), and because the one new deal in marketing (Blue Owl's $513M) is hyperscaler-leased. It is held at 38 rather than lower because the two quoted deals are the test itself: Compass is a multi-tenant developer, so a CMPDC print at or above $300M would likely qualify on the tenant table, and a deal that is quoted for four weeks is more often repriced than pulled; because the last data-centre ABS to print (EdgeCore's $521M single-A at T+325 on October 20) was a non-hyperscaler operator's paper and cleared at a spread Saturn's A- at T+345 has since matched, so the rated bid for operator paper exists; because the 2026 market printed $21.9B year-to-date with multi-tenant and neocloud-leased pools in it before the drought; because a December-February window is three months and the conduit calendar restarts in January; and because high-yield spreads at 352 bp are a repricing rather than a shutdown for liquid credit. It is held below 50 because the tenant-table test may exclude Compass if its pool is majority hyperscaler-leased, because the MSC deal's tenancy is not in the packet, and because a four-week drought on a quoted deal is itself evidence the price has not been found. A hit says the bifurcation was a four-week calendar effect; a miss through February says the market has split by signature for at least a quarter, which is the big story's claim.

PREDICTION 02 SOFTWARE

28%

BY JANUARY 31, 2027

By January 31, 2027 Artificial Analysis publishes an Intelligence Index v5 score of 50 or higher for a model whose weights are downloadable from Hugging Face without an application gate under a licence that permits commercial use (any territory restriction or revenue threshold noted but not disqualifying), that is, the open floor closes the gap to the closed leader (65) to 15 points or less from 19 within the fourth hypothesis's second measured quarter.

TRIGGER: Hit if, by the deadline, artificialanalysis.ai lists a model at 50 or higher on Intelligence Index v5 (any effort setting) and the model's Hugging Face repository contains weight files and a licence file permitting commercial use (Apache 2.0, MIT, a Qwen or Tencent community licence with thresholds, or similar) with no approval gate; a score of 47-49 on such a model, or a score of 50 or higher on a model whose weights are gated or published only on ModelScope, is a partial; no downloadable model at 47 or higher by the deadline is a miss; if Artificial Analysis rebases the index the resolver maps the threshold to the equivalent position relative to the closed leader (15 points behind or closer). The bet sits well below even because the best downloadable score has been Kimi K3.1 Pro's 46 since November 4 with nothing closer than Hy4 at 35 landing since, because Alibaba shipped Qwen4-Max closed and its open Qwen4 models are 'coming soon' with no date, because Thinking Machines' two open releases score 31 and 20.1, and because the closed leader moved five points in the last three weeks and may move again. It is held above 20 because the week delivered more open artifacts than any in the series (Hy4 at 770B, Hy4 flash, Inkling 2 Mini, Aion), because a Qwen4 mid-size open model or a Kimi K3.2 is the kind of release that has moved the floor by 5-10 points before (Kimi K3.1 Pro took it from ~40 to 46), and because the fourth hypothesis's refutation clock (gap above ten for two quarters, ending late April) gives Chinese labs an incentive this issue cannot measure. A hit moves the fourth hypothesis toward supported before the second quarter closes; a miss is the third consecutive month in which the open floor rose in artifacts and not in score.

PREDICTION 03 HARDWARE

61%

BY FEBRUARY 28, 2027

By February 28, 2027 a TrendForce report, or a memory supplier's earnings release or call (SK hynix, Samsung, Micron), discloses or confirms a 2027 HBM4 contract price at or above $3.00 per gigabit (the bottom of the roadshow range; +46% on the 2026 $2.05), that is, the 2027 HBM4 price step lands at or above 1.5x, the opposite side of the W40 bet that the settled price lands below the suppliers' ask.

TRIGGER: Hit if, by the deadline, a TrendForce press release or research bulletin on trendforce.com states a settled or contracted 2027 HBM4 price (any customer, any stack height) of $3.00 per gigabit or more, or a supplier's earnings materials or call transcript (SK hynix Q4 2026 in late January, Samsung Q4 in late January, Micron FQ1 2027 in December or FQ2 in March) states an HBM4 price increase for 2027 of 46% or more over 2026 or a per-gigabit figure at or above $3.00; a disclosed or reported contract price of $2.60-2.99 per gigabit (+27% to +46%), or a supplier statement of a 2027 HBM4 price increase of 25-45% without a per-gigabit figure, is a partial; a disclosed price below $2.60, a supplier describing 2027 HBM4 pricing as flat to modestly up, or no disclosure by any of the named sources by the deadline is a miss (the resolver notes that TrendForce's per-gigabit figures are blended estimates and uses the lowest customer-specific figure it publishes if it publishes more than one; a TrendForce note that merely reiterates the $3.00-5.50 forecast range without stating a settled or contracted price does not count as a hit). The bet sits above even because TrendForce put the whole 2027 range above $3.00, because Nvidia's supply commitments rose $111B in a quarter with multi-year memory procurement the largest component (a buyer committing at the top of a range does not then negotiate it down), because HBM is forecast to take 30% of DRAM wafers with bits up 65.7% while PC DRAM bits fall, and because Korean press reported SK hynix's HBM4 at a mid-$500s 12-high stack price with Samsung seeking parity (W46, grade 4). It is held at 61 rather than higher because the W40 house bet (the settled price lands below the suppliers' ask) is the opposing position and is still open, because the trigger requires a disclosure that suppliers have historically avoided (Micron and SK hynix describe HBM pricing qualitatively), because the Taiwan slip reports, if the HBM4 leg is true, could mean customers delay rather than pay, and because a $3.00 floor on a range that runs to $5.50 is the easiest leg of the roadshow to confirm and the hardest to contradict with a qualitative statement. A hit confirms the fifth hypothesis's components-bind-through-2027 leg with a price; a miss says the buyers' order books bought allocation without paying the step.

PREDICTION 04 NETWORKING

64%

BY FEBRUARY 28, 2027

By February 28, 2027 Nvidia's Q4 FY27 results (expected in the last week of February) show networking revenue growing faster year on year than data-centre compute revenue for a second consecutive quarter (Q3 FY27: networking above $21B from $8.2B, about +156%, against compute of roughly $86.6B from $43.0B, about +101%, house arithmetic on the disclosed segments), that is, interconnect outgrows compute at the largest vendor for two quarters running, which is the third hypothesis's test applied at the vendor layer where the operator-layer leg failed this week.

TRIGGER: Hit if, by the deadline, Nvidia's Q4 FY27 press release, CFO commentary, 10-K or the prepared remarks on its earnings call (the transcript on investor.nvidia.com) discloses networking revenue and data-centre revenue such that networking's year-on-year growth rate exceeds the year-on-year growth rate of data-centre revenue less networking by 10 percentage points or more (the resolver computes both from the disclosed figures and the prior-year Q4 FY26 figures in the same filing or the FY26 10-K); networking growing faster than compute by less than 10 percentage points is a partial; networking growing slower than compute, or Nvidia ceasing to disclose networking separately, is a miss (the resolver records the disclosure change). If Nvidia's Q4 FY27 release falls after February 28, the resolver uses the first release after the deadline, records the timing, and scores on the same terms. The bet sits above even because networking grew about 2.6x against compute's about 2x (+101%) in Q3 on a Blackwell-to-Rubin transition in which the NVL72 rack carries more switch and optics content per GPU than any prior platform, because Nvidia said CPO scale-out is in volume production and Spectrum-X and Quantum-X ship inside every Rubin rack (Dell's integrated networking rack to Nscale is the physical form), because the ~$20B of first-quarter Rubin revenue implies the networking attach rate rises with the ramp, and because Q4's $130B guide is 12% above Q3 with Rubin a larger share. It is held at 64 rather than higher because a single quarter's growth differential on a transition can reverse when compute catches up (Q4 FY26's networking base of about $11B is higher than Q3 FY26's $8.2B, so the comparison gets harder by about 30 points of growth), because Nvidia's disclosure of the split is a press-release convention it can change, and because the margin guide to 71-72% says the Rubin ramp's cost is landing in compute, which could depress compute growth and flatter the comparison for a reason that has nothing to do with interconnect value. A hit gives the third hypothesis its first two-quarter leg anywhere in the stack; a miss leaves it strained on both layers.

PREDICTION 05 POWER

76%

BY DECEMBER 31, 2026

By December 31, 2026 PJM posts results for the 2029/2030 Base Residual Auction (scheduled December 9-15) showing the RTO-wide clearing price at the $337.83/MW-day cap, with the auction held on schedule rather than postponed on the seven-state advocates' request, that is, the Independent Market Monitor's projection holds and the capacity market prices the data-centre forecast at its ceiling for a third consecutive auction.

TRIGGER: Hit if, by the deadline, PJM's posted 2029/2030 BRA results (the results report on pjm.com) show the RTO resource clearing price equal to the cap of $337.83/MW-day (within $0.01) and the auction was conducted within the December 9-15 window or within seven days of it; an RTO clearing price within 5% of the cap but below it, or an on-schedule auction with the cap binding only in constrained zones (Dominion, BGE, etc.) and the RTO price below it, is a partial; a postponement of the auction past December 31 on the PJM Board's or FERC's action, or an RTO clearing price more than 5% below the cap, is a miss. The bet sits well above even because the Monitor, who has the data, projects a clear at the cap, because the last two auctions cleared at or near their caps with data-centre load the stated driver, because PJM's Board has rejected postponement requests before and FERC has no pending order to delay, and because the Monitor's own 4,017 MW of cancelled or delayed load has already been argued into the ER27-454 record without changing the auction parameters. It is held at 76 rather than higher because seven states' advocates asking the Board to seek a FERC postponement two weeks before the auction is an unusual and specific request, because a Board that is also defending the RBP at FERC has an incentive to show flexibility, because the Monitor's shortfall arithmetic (3,055 MW adjusted against PJM's 6,831 MW) could be partly reflected in offers, and because a clear within a few dollars of the cap rather than at it is a real possibility that the trigger scores as a partial. A hit records the capacity market pricing the 2028-30 data-centre forecast at its ceiling and adds $5-10B (the Monitor's estimate) to the cost allocation the fifth hypothesis says precedes the physical constraint; a miss by postponement says the advocates' carve-out argument won before the auction ran.

PREDICTION 06 CAPITAL

57%

BY DECEMBER 31, 2026

By December 31, 2026 Anthropic's Class A shares (ANT) close below the $83.00 IPO price on at least one Nasdaq trading day, that is, the public read on the labs' commitment schedules that the W46 synthesis named trades through the price the lab's own suppliers paid for their Class N stake within six weeks of listing.

TRIGGER: Hit if, by the deadline, Nasdaq's official closing price for ANT on any trading day is below $83.00 (an intraday low below $83 with a close at or above it does not count; November 20's $83.05 low with an $88.91 close is the baseline); a close between $83.00 and $85.00 inclusive on any day with no close below $83 is a partial; no close below $85 by the deadline is a miss. The resolver uses Nasdaq's historical data and ignores any stock split or distribution. The bet sits above even because the stock gave back the entire pop in two sessions and touched $83.05, because the Nasdaq fell 10% in the listing week with a Federal Reserve governor floating a 50 bp December hike (odds about 98% for a move at the December 8-9 meeting), because the lock-up structure and the 933.7M Class N shares reserved for the Donation Matching Plan (roughly $100B at day-one prices; counterbrief, outside the graded packet) are overhangs the market has had two days to price, and because six weeks is long enough for one bad tape day on a stock that trades at about 15x run-rate revenue. It is held below 65 because the offshore perpetual-futures reference near $95 (Blockworks, outside the graded packet) and the $83.05 bounce suggest the issue price is a level buyers defend, because seven suppliers hold $26B at $83 and underwriters have a stabilisation bid for the first weeks, because the 10-year Treasury rallied 23 bp in the same week, and because the stock closing above issue through December would be the ordinary outcome for an IPO this oversubscribed. A hit says the market is pricing the lab's roughly $895B of commitments below the terms its suppliers paid, which is the chain the W46 connection asked to be stressed; a miss says the suppliers' price held as a floor.

TRACK RECORD

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

Predictions made	168	Resolved (hit / partial / miss)	89 (41 / 20 / 28)
Hit rate (partial = half)	57%	Brier score (0 = perfect)	0.194

Bold (<55%): 9 resolved, 67% hit rate vs 41% mean confidence · Core (55-80%): 75 resolved, 54% hit rate vs 65% mean confidence · High-conviction (>80%): 5 resolved, 90% 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 25

WATCH 01

Qwen4-Max Preview retires; any Qwen4 open-weight repository on Hugging Face

Alibaba's production launch said smaller open Qwen4 models are 'coming soon'; the preview's end is the first date on which silence becomes a datapoint for p159 (36%) and p164 (28%), and a Qwen4 repository with weights and a licence would move both and the fourth hypothesis's gap.

NOVEMBER 30

WATCH 02

PJM ER27-454 comment deadline; CREFC's fifth weekly read on data-centre ABS/CMBS; AWS re:Invent opens

The RBP record closes with the Monitor's 4,017 MW and the seven-state postponement letter in it; a fifth week without a data-centre securitisation pricing extends the bifurcation the big story claims (p163); re:Invent is the first hyperscaler stage since Nvidia's print for Trainium and rack pricing.

DECEMBER 1

WATCH 03

Gemini 4 Argon and Carbon step from $2 / $10 to $4 / $20; Carbon's $0.10 cache rate ends

The first list increase inside the mid-tier band in the series resolves two levers (mid-tier band; Opus cache threshold) and the non-list-meter pattern's confirming leg; an extension or cancellation breaks the pattern and flips the cache lever.

DECEMBER 8-9

WATCH 04

FOMC decision (a 50 bp hike floated by Governor Waller, ~98% odds of a move); Broward County's 50 MW moratorium vote (December 8); PJM BRA opens December 9

A hike into a week when the 10-year already rallied to 4.92% resets the discount rate on every lease the house measurement priced at 5.95% and bears on p168 (ANT below $83); Broward is the next county on the moratorium map; the BRA is p167's event.

DECEMBER 9-15, RESULTS BY ABOUT DECEMBER 23

WATCH 05

PJM 2029/2030 Base Residual Auction

p167 (76%) is written on an on-schedule auction clearing at the $337.83/MW-day cap; a postponement on the advocates' request or a clear below the cap would be the first time since 2024 the data-centre forecast did not price at the ceiling, and the Monitor's $5-10B estimate is the cost allocation the fifth hypothesis says precedes the constraint.

DECEMBER 31, THEN JANUARY 10 AND JANUARY 14

WATCH 06

NDAA conference (Chip Security, AI OVERWATCH, MATCH riders); US and China tariff and export-control suspension cliffs (January 10); OpenAI dots billing begins (January 14)

The NDAA riders are the legislative half of the two-layer export regime (BIS remote-access rule plus anti-distillation) the second-order section describes; the January 10 cliffs are the dated test of the Shenzhen non-deal; dots billing at $50-500 per agent per month is the first per-agent subscription meter to turn on by default (The Application Layer).

LATE JANUARY AND LATE FEBRUARY 2027

WATCH 07

SK hynix and Samsung Q4 calls (HBM4 2027 pricing; p165); Nvidia Q4 FY27 (networking versus compute; p166); Micron FQ1 in December as the first supplier read

Three supplier prints decide two of this week's bets: a per-gigabit or percentage HBM4 disclosure at or above $3.00 / +46% confirms the components-bind leg with a price, and Nvidia's segment split decides whether interconnect outgrew compute for a second quarter.

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 31

Week 47 of 2026 · November 21, 2026

WEB

brianletort.ai/industry

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