EDITION 053

SUNDAY, NOVEMBER 22, 2026

MIND CLUB

AI & Technology

Intelligence

AI OBSERVATORY · PDMC.AI

SPECIAL COVERAGE · THE EDITION IN FOUR SQUARES

03S · P. 09

Nvidia's leadership changed. The spending kept climbing.

REG

01 · P. 04

A Riemann proof passes Lean; its author credits Claude

DATA

02 · P. 04

Nvidia ends its worst week since 2008, as reported

DATA

03 · P. 05

OpenRouter drops three inference providers over prompt-data incidents

CO

WHERE THIS COULD BE WRONG · P. 04

A proof nobody can yet read is a strange kind of result.

THIS EDITION

Riemann proof passes Lean; author credits Claude

BUSINESS

A RIEMANN PROOF PASSES LEAN'S CHECKS, AND NVIDIA ENDS THE WEEK DOWN 22.92%

The weekend's biggest AI story came from mathematics: the Lean community says a machine-checked proof of the Riemann Hypothesis is correct; its author, a mathematician at Anthropic, says he made it with Claude, though no one has yet explained it in words. Markets were shut after the week's sell-off; Nvidia had its worst week since 2008, as reported, and has not answered a report of an SEC inquiry. This edition also carries Special Coverage on chips and on Anthropic and SpaceX.

QQQ · FRI 11.20	SMH (SEMIS)	NVDA	AMD
608.45	**400.06**	**125.53**	**424.39**
WEEK -10.51% FRIDAY CLOSE	WEEK -17.81% FRIDAY CLOSE	WEEK -22.92% FRI -14.01%	WEEK -20.54% FRI -13.25%
ANT	**10Y TREASURY**	**AI DEALS ≥$100M**	**HYPERSCALER CAPEX**
88.91	**4.92%**	**$37.3B**	**$725B+**
$83 IPO PRICE FRI -8.50%	WEEK ABOUT -23 BP 11.20 CLOSE	THIS WEEK, 11 DEALS $26.0B IS ONE PLACEMENT	2026 GUIDE, BIG 4 NOT ADDED TO TOTALS

VERIFICATION AT A GLANCE

CLAIM CHECK	CALENDAR DATES	SOURCE WATCH	OPEN CALLS	CORRECTIONS
8	**3/9**	**30**	**8**	**5**
2 match · 3 partial · 3 open	confirmed at the source	documents · not re-read this edition	2 new · 5 old results printed	in the log · 2 new this week

#	FIVE THINGS FOR THE WEEKEND	SOURCES	PRIORITY
01	**A Riemann proof passes Lean; its author credits Claude.** The Lean and Mathlib communities say the statement and the proof are correct; the ideas behind it have yet to be explained.	DATA PRESS	MUST READ
02	**Nvidia ends its worst week since 2008, as reported.** It closed at $125.53, down 22.92% for the week; it has not answered a report of an SEC inquiry, and its 10-Q isn't filed yet.	DATA REG PRESS	MUST READ
03	**OpenRouter drops three inference providers over prompt-data incidents.** It also pulled 125 endpoints from its zero-retention list and says it was not breached itself.	CO	MUST READ
04	**Thinking Machines posts an open-weights model with a 20 GB checkpoint.** Its own page gives two counts of active parameters.	CO	KNOW IT
05	**Bright spot: OpenAI publishes a chemistry test with its limits attached.** Its best model scored 84.5%, but a top model earned about half the credit on choosing compounds.	CO	SCAN

SPECIAL COVERAGE IN THIS EDITION

Special Coverage: the AI chip cycle, with the money loop and what the build-out means for the economy. 03S.
Special Coverage: Anthropic and SpaceX on the same quarters, and the compute deal read two ways. 03T.

HOW TO READ THIS BRIEF

MUST READ understand this weekend. KNOW IT secondary. SCAN awareness only. VERIFIED means a primary source was checked, CLAIM means a company or press statement, ANALYSIS means our interpretation, FORECAST means our odds, and SPECULATION means plausible but unverified. The red boxes are the **red team**: the strongest case against the implication just drawn.

Source badges show what each item rests on: REG a registry or filing GOV agency or government PEER peer-reviewed journal PRE preprint, not peer-reviewed CO company release, model card or pricing page DATA market data, an independent benchmark or a registry search PRESS news report.

Claim Check verdicts compare a headline with its primary document: MATCH PARTIAL MISMATCH OPEN (no primary document found yet). A verdict says what differs, never why.

READING PRIORITY INDEX

MUST READ	The Riemann proof · Nvidia's week · OpenRouter's providers · Special Coverage 03S and 03T · Where this brief could be wrong
KNOW IT	Inkling 2 Mini · The week in review · AI Theme Pulse · Compute & Capital Map · Benchmark reality check · Try this weekend · The A to S on Intelligence · Public Calls Ledger · The week ahead
SCAN	MedChemBench · Paper of the week · The decoder · Outlook, condensed

IN THIS EDITION

00	At a glance	02	09	The decoder	25
01	Five things for the weekend	04	10	Tracking Super Intelligence (SI)	28
02	Five talking points	06	11	Pattern recognition	30
03	The week in review	07	12	What may be underestimated	31
03S	Special Coverage: The AI Chip Cycle	09	13	Where this brief could be wrong	31
03T	Special Coverage: Anthropic and SpaceX, Like for Like	17	14	Outlook, condensed	31
04	AI Theme Pulse	20	15	The week ahead	32
05	Compute & Capital Map	21	16	Public Calls Ledger	32
06	Benchmark reality check	22	17	Ask Mind Club and corrections	34
07	Try this weekend	24	18	Sources	35
08	Paper of the week	25			

MIND CLUB IS PUBLISHED BY TORPEDO PUBLISHING LLC

FIVE THINGS FOR THE WEEKEND

Markets closed before the window opened, so items 01 and 03 are the weekend's own news. Item 02 adds Friday's official closes and what has and hasn't been filed since. Items 04 and 05 are dated Friday 11.20, from posts that may predate the window; we carry them with their dates.

01 A Riemann proof passes Lean; its author credits Claude MUST READ

WHAT HAPPENED

On Friday 11.20, Levent Alpöge, a mathematician on Anthropic's technical staff, posted a short note and a Lean library of about 17 million lines claiming a proof of the Riemann Hypothesis, and of the Generalized Riemann Hypothesis for Dirichlet L-functions. He says he found it working with Claude. CLAIM

On Saturday 11.21, the Lean and Mathlib communities, endorsed by the Lean FRO and the Mathlib Initiative, posted a statement of full confidence that "both the statement and the proof of the Riemann Hypothesis are correct." Their basis: the theorem matches Mathlib's own vetted statement of the hypothesis, and several independent proof checkers accepted the exported proof. MIT's Andrew Sutherland reports that all 175,187 modules compiled with no errors. The statement says it looks forward to work that "will help explain the new ideas in this proof." VERIFIED

WHY IT MATTERS

If it holds, this is one of the most famous open problems in mathematics settled with an AI model in the loop. A computer check of this kind is strong evidence that the logic holds; it is not yet the understanding mathematicians expect from a proof, and the formal result rests on Mathlib's definitions being right. ANALYSIS

WHAT WE'RE WATCHING

The remaining checker runs, a human-readable account of the ideas, and whether a journal takes it on. The prize rules of the Clay Mathematics Institute require publication and a waiting period.

CLAIM CHECK MATCH The Lean community's post and Sutherland's verification page say what we print. OPEN How much of the work was Claude's and how much Alpöge's rests on his own account and press reports, and Anthropic has not issued a statement.

RED TEAM

A proof nobody can yet read is a strange kind of result. A written version would reportedly run past 250,000 pages, so it may be years before people agree on what was shown, and why.

SOURCE · LEAN COMMUNITY STATEMENT (11.21) DATA · SUTHERLAND, VERIFICATION PAGE DATA · ALPÖGE, NOTE PRE · OFFICECHAI (11.21) PRESS

02 Nvidia ends its worst week since 2008, as reported MUST READ

WHAT HAPPENED

Nvidia closed Friday 11.20 at $125.53, down 14.01% on the day and 22.92% for the week from $162.85 on 11.13. AMD fell 20.54% for the week and SMH 17.81%. VERIFIED Investing.com calls it Nvidia's worst week since 2008 and its seventh straight weekly fall. CLAIM

Reuters reported on 11.20, citing two people briefed on the matter, that SEC enforcement staff are examining "potential accounting issues" at Nvidia; The Wall Street Journal first reported an SEC inquiry on 11.17. Reuters could not say when it began or what it covers; the SEC declined to comment. CLAIM Through early Sunday we found no statement from

Nvidia and no filing on it. EDGAR shows no Nvidia 8-K after the 11.17 results, and no 10-Q for the quarter yet; the CFO commentary filed with the results says Nvidia plans to file the 10-Q "in the coming weeks." VERIFIED

WHY IT MATTERS

The results on 11.17 showed demand still climbing, with a $130 billion guide. The week's fall suggests investors are pricing how the company is run and reported, not only how many chips it sells. The 10-Q may be the first filing to speak to either. ANALYSIS

WHAT WE'RE WATCHING

The 10-Q, normally due about 12.04 (our call #29, 75%), any company statement on the reported inquiry, and the chief executive search (call #30, 25%).

CLAIM CHECK MATCH Friday's close and the weekly change match the daily bars; EDGAR's filing list shows what we print. OPEN The SEC inquiry rests on Reuters and the Journal; no filing confirms it.

RED TEAM

An inquiry reported by unnamed people is not a finding, and many SEC inquiries end with nothing. A quiet company is normal while one is open.

SOURCE · SEC EDGAR, NVIDIA FILINGS REG · NVIDIA, Q3 FY2027 RESULTS (11.17) CO · ROBINHOOD DAILY BARS DATA · REUTERS, VIA KOREA TIMES (11.21) PRESS · INVESTING.COM (11.20) PRESS

03 OpenRouter drops three inference providers over prompt-data incidents MUST READ

WHAT HAPPENED

OpenRouter, a service that routes requests to many model providers, said on Sunday 11.22 that it has removed three providers, which it does not name, and pulled 125 endpoints from its Zero Data Retention list, after research credited to the security firm Dragos. The post links incident notices from Together (prompt leakage for requests in June and July 2026), Phala (a preliminary incident report in November) and Novita (an announcement on X). OpenRouter says it was not breached itself. VERIFIED Dragos says it sent about two million test prompts carrying planted secrets to AI providers over July to November, and that the largest of three breaches it found involved "a couple hundred" of its planted cloud keys. CLAIM

WHY IT MATTERS

A routing service sends your prompt to whichever provider is cheapest or fastest. That means your data's safety rests on companies you may never have picked, and a "zero retention" label is only as good as the provider behind it. ANALYSIS

WHAT WE'RE WATCHING

Statements from the three providers, whether Dragos names them, and whether other routers publish their provider checks. Idea B in 07 Try this weekend shows how to check your own settings.

CLAIM CHECK PARTIAL OpenRouter's post supports the removals and the 125 endpoints; it doesn't say which three providers were removed. OPEN The Dragos figures are Dragos's own.

RED TEAM

The incidents vary in kind, and one rests on a social post. Removing providers may say more about OpenRouter's caution than about how widespread the problem is.

SOURCE · OPENROUTER (11.22) CO

04 Thinking Machines posts an open-weights model with a 20 GB checkpoint KNOW IT

WHAT HAPPENED

Thinking Machines released Inkling 2 Mini on Friday 11.20, a “looped” mixture-of-experts model with 26.4 billion total parameters and a checkpoint of about 20 GB; its model card recommends 32 GB or more of memory to run it. Its page reports 79.2% on SWE-bench Verified (bash only) and 84.5% and 49.3% on ARC-AGI versions 1 and 2, and notes that “evaluation setups vary by model.” The page gives 3.4 billion active parameters, or 5.4 billion counting the looped layers; the release page names no license, and the Hugging Face model card lists Apache 2.0. VERIFIED

WHY IT MATTERS

A model whose card recommends 32 GB of memory, one high-end graphics card or a 32 GB Apple laptop, can be run in house, keeping data off other people’s servers, which is the risk item 03 is about. ANALYSIS

CLAIM CHECK PARTIAL The scores are on the page as printed; the page uses two active-parameter counts. The release page states no license; the Hugging Face model card lists Apache 2.0.

RED TEAM

Self-reported scores on setups that “vary by model” don’t compare cleanly with anyone else’s. Check the license on the model card yourself before using it at work.

SOURCE · THINKING MACHINES (11.20) CO

05 Bright spot: OpenAI publishes a chemistry test with its limits attached SCAN

WHAT HAPPENED

OpenAI introduced MedChemBench on Friday 11.20: 99 medicinal chemistry problems drawn from 117 papers and patents, covering 761 molecular structures. GPT-6.2 Astra scored 84.5%, GPT-6 Astra 83.3% and GPT-5.6 Sol 58.2%. The top models did well at reconstructing structures but earned only about half the credit on selecting compounds and predicting properties. OpenAI says the models may have seen the source papers, that each answer was a single try without tools, and that the test does not show the models can discover drugs. VERIFIED

WHY IT MATTERS

A vendor that prints its own benchmark’s weak spots gives chemists something to test against, rather than a headline to argue with. ANALYSIS

CLAIM CHECK PARTIAL The scores and caveats are as printed, at the highest reasoning setting; the test is OpenAI’s own, and the page counts eleven or nine models tested but shows five full scores.

SOURCE · OPENAI, MEDCHEMBENCH (11.20) CO

Also in the window. Bloomberg reports Cognition is in talks to raise $2 billion or more at about $75 billion, and the Journal and Reuters report Databricks is in early talks at about $225 billion, as relayed by AI Funding Tracker PRESS. Researchers at NebuSec say flaws in Meta’s Muse agent let a webpage run code inside it and quietly send out users’ chats and saved memories, and that Meta patched them before publication; a Cloud Security Alliance note relays this and cites no Meta statement PRESS. Talks at the United Nations on autonomous weapons ended with states agreeing only to “elaborate” an accord, as reported PRESS.

02 / FIVE TALKING POINTS THIS WEEK’S CONCLUSIONS

FIVE TALKING POINTS

Five conclusions from the week. One bears on biotech and pharma, flagged.

01 **The sell-off hit the chips, not the demand data.** Nvidia guided to $130 billion on 11.17 and then fell 22.92% in the week, a week that also brought a reported SEC inquiry. On the record so far, the order books still read strong. ANALYSIS

02 **The biggest IPO held its price, barely.** ANT ended Friday at $88.91, about 7% above its $83 offer, in a week the S&P 500 closed in correction. ANALYSIS

03 **A checked proof moves the question from "is it right?" to "what does it mean?"** The Lean community backs the Riemann proof; people still have to explain it. ANALYSIS

04 **The provider behind the model is now a security question.** OpenRouter's removals follow last week's provider breaches. Ask who runs your model, not just which model it is. ANALYSIS

05 **Isomorphic raised $5 billion and signed Merck, but nothing is registered.** BIOTECH & PHARMA A ClinicalTrials.gov search by sponsor returns no Isomorphic trial, with about five weeks left in the year. VERIFIED

THE WEEK IN REVIEW

The week's top items from the weekday editions, as printed after their fact-checks, with Friday's official closes. The full Special Coverage desk in 03S replaces the short version that ran on Wednesday 11.18 (Edition 050).

DAY	WHAT HAPPENED	LABEL
Mon 11.16	Chips slide the day before Nvidia reports, SOXX down 5.43%; Commerce releases its remote-access chip rule; California sues OpenAI, alleging it lost control of its agents	Verified and Claim
Tue 11.17	Nvidia reports $115.8 billion and guides to $130 billion; GPT-6.2 Astra and Sol arrive; Microsoft ships its own models; Isomorphic Labs raises $5 billion	Verified
Wed 11.18	Nvidia falls 2.96% after results and its chair reportedly rules out Huang's return; Palo Alto agrees to buy Applied Compute for $4.5 billion in stock; Anthropic prices after the close	Verified and Claim
Thu 11.19	Anthropic opens at $111.50 and closes 17.1% above its $83 offer, as reported; the remote-access rule is published and in force; Gemini 4 Neon launches; Merck hires Isomorphic	Verified and Claim
Fri 11.20	The S&P 500 closes in correction; Nvidia falls 14.01%; Reuters reports an SEC inquiry into potential accounting issues at Nvidia; a House committee releases China chip license data	Verified and Claim

FRIDAY 11.20 CLOSE	CLOSE	WEEK
QQQ	608.45	-10.51%
SMH	400.06	-17.81%
NVDA	125.53	-22.92%
AMD	424.39	-20.54%
TSM	318.71	-20.07%

FRIDAY 11.20 CLOSE	CLOSE	WEEK
ORCL	101.19	-19.27%
PLTR	143.49	-14.58%
META	560.84	-12.60%
AVGO	304.41	-12.28%
GOOGL	265.29	-10.31%
MSFT	484.52	-7.51%
AMZN	246.82	-7.44%
ANT	88.91	listed 11.19

Regular-session daily bar closes from Robinhood; changes computed against the 11.13 closes. The 10-year Treasury yield closed at 4.92% on 11.20, as reported, about 23 basis points below the 5.15% of 11.13 (Treasury), a move toward safety. DATA PRESS

Influence map. The week's center of gravity moved from the labs to the people who check them: an exchange's first auction, a regulator's reported inquiry, a rule in the Federal Register and, on Saturday, a community of proof checkers. ANALYSIS

RED TEAM

A fall this sharp after record results has happened before and reversed. One week of selling says what investors fear, not what the businesses will report.

NVIDIA'S LEADERSHIP CHANGED. THE SPENDING KEPT CLIMBING.

Six years of filings show AI chip demand still rising with the budgets behind it. The open questions are who pays for the next round, how much of Nvidia's deal-making left with its founder, and how much of the U.S. economy now rides on the answer.

Sources: company 10-K, 10-Q, 8-K, 6-K and S-1 filings; Bureau of Economic Analysis and Federal Reserve data; "reported" marks press accounts. This is the full version of the desk that ran in short form on Wednesday (Edition 050); our coverage of Nvidia's board review ran in Editions 042 and 044. Mind Club takes no position on any company or security. Mind Club research is developed through a human-led AI workflow using specialized expert agents, cross-model analytical challenge, primary-source review and substantial editorial direction. Independent analytical perspectives are used to test assumptions, identify weaknesses and refine conclusions. The human editor retains responsibility for the final work product. Some AI companies covered here supply models used in that workflow.

THE SHORT VERSION

The structure is intact, and on the record so far, demand is still climbing. Microsoft, Alphabet, Amazon and Meta spent $202 billion on property and equipment in the September quarter. Every supplier layer we checked posted record AI-linked sales, and Nvidia's data center sales now run at about half of what the four biggest buyers spend on property and equipment. Huang's exit matters in three places: Nvidia's role as lender and guarantor to its own customers, its standing with governments on exports, and the bets beyond its current chips. It is not, so far, showing up in orders. The warning signs that came before past chip busts are now in financing, not demand, and the build-out is now big enough to move the U.S. economy. ANALYSIS

What happened

Nvidia's board accepted Jen-Hsun Huang's resignation as president, chief executive and director on 11.11 and announced it the next day. Colette Kress, the finance chief since 2013, is interim chief executive and keeps the finance job; Harvey C. Jones, a director since 1993, is the company's first chair. Nvidia has given no reason for the resignation. VERIFIED Five days later the company reported data center revenue of $107.6 billion, up 110%, and guided to $130.0 billion of total revenue, plus or minus 3%, for a quarter with an extra week, with a gross margin of 71% to 72%, down from 74.8%. The guide assumes no data center compute sales to China. VERIFIED The board chair reportedly said the board does not intend to reinstate Huang, and Huang was reported to have said he would be "honored" to return. CLAIM

Since Wednesday. Nvidia closed at $125.53 on Friday 11.20, down 14.01% on the day and 22.92% for the week from $162.85 on 11.13 (Robinhood data). VERIFIED Reuters (11.20) reported that SEC enforcement staff are examining potential accounting issues at Nvidia, after The Wall Street Journal (11.17) reported an SEC inquiry into its disclosures; no filing confirms either, and we found no response from Nvidia by Sunday morning. CLAIM On the chief executive search, no change: no 8-K since 11.17, and Colette Kress remains interim chief executive. Commerce's remote-access rule was published and took effect on 11.19 (FR Doc 2026-23529): it requires licenses for advanced chips bound for data centers that provide more than 12 million TPP of aggregate computing power to or for companies headquartered, or with an ultimate parent headquartered, in Macau or a Country Group D:5 destination, which include China; adds license requirements for Indonesia, Malaysia, Singapore, Thailand and Brazil, with an exclusion for U.S.-headquartered buyers who file monthly reports; formally rescinds the AI Diffusion Rule; and takes comments to 01.19.2027. Anthropic's shares, which the placement buyers agreed to buy at the $83 IPO price, closed at $97.17 on 11.19, their first day, and $88.91 on 11.20, about 7% above that price. VERIFIED

Those numbers answer the narrow question of whether customers walked. They did not. The wider question is whether the spending behind them can last, and that needs a longer look. ANALYSIS

Follow the money down the stack, 2020 to now

Demand for AI chips starts with a handful of buyers and runs down through every layer that serves them. Two things stand out. The supplier layers grew far faster than the spending above them: the Big Four's capital spending roughly quadrupled from 2020 to 2025, while Nvidia's data center revenue grew about 29 times over the matching fiscal years. And the one year spending fell, 2023, was a boom year for Nvidia: Amazon cut capital spending and Meta trimmed it, while Nvidia's data center sales more than tripled. A flat budget can still be a boom for the supplier whose share of it is rising. ANALYSIS

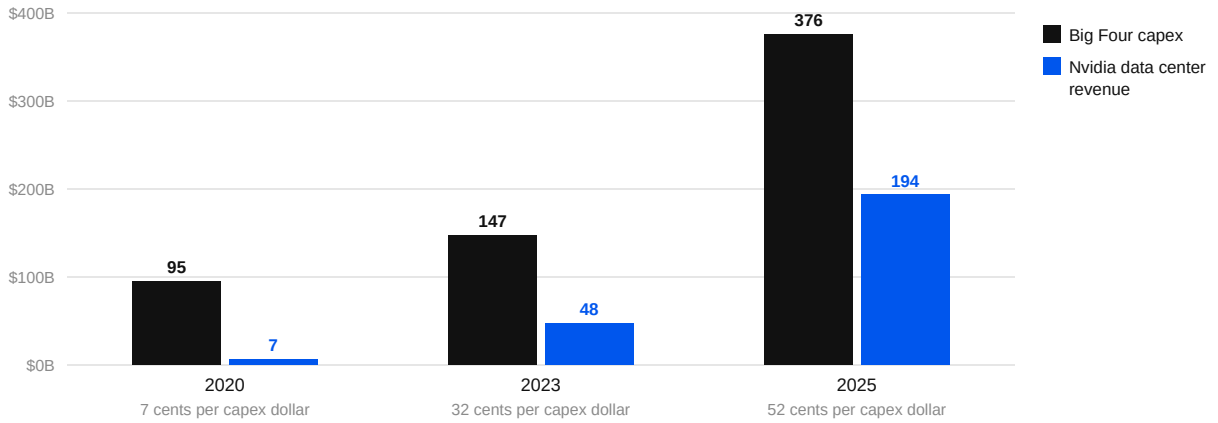

$B. Big Four: Microsoft, Alphabet, Amazon, Meta, as each reports capital spending. Nvidia: fiscal years ending late January of the next year.

LAYER	COMPANY AND MEASURE, $B	2020	2023	2025	LATEST QUARTER	VS YEAR AGO
Who spends	Microsoft, capital spending	17.6	35.2	83.1	39.4	+103%
	Alphabet	22.3	32.3	91.4	51.4	+114%
	Amazon	40.1	52.7	131.8	77.3	+120%
	Meta	15.2	27.0	69.7	33.9	+80%
	Oracle, quarter to August	1.6	8.7	21.2	28.5	+235%
Who rents out the machines	Amazon Web Services, revenue	45.4	90.8	128.7	46.0	+39%
	Google Cloud, revenue	13.1	33.1	58.7	30.7	+103%
	CoreWeave, revenue	not reported	0.2	5.1	3.7	+168%
Who uses them	Anthropic, revenue	not reported	0.04	4.6	17.0	+1,127%
Who designs the chips	Nvidia, data center revenue (fiscal years)	6.7	47.5	193.7	107.6	+110%
	AMD, data center revenue (2020 recast)	1.7	6.5	16.6	8.5	+97%
	Broadcom, AI chips (2025 reported)	not disclosed	not disclosed	about 20	16.7	+221%

LAYER	COMPANY AND MEASURE, $B	2020	2023	2025	LATEST QUARTER	VS YEAR AGO
Who makes them	TSMC, revenue	45.5	69.3	122.4	46.7	+41%
	SK hynix, revenue, trillion won	31.9	32.8	97.1	104.6	+328%
	Micron, revenue (fiscal years to early September)	21.4	15.5	37.4	54.2	+379%
The whole market	World chip sales	440	527	796	about 1,500 (2026 forecast)	

Latest quarter is the September 2026 calendar quarter unless marked; Nvidia's ended 10.25, Broadcom's 08.02, Micron's is fiscal Q4 2026. Growth against the same quarter a year earlier. Every company figure from its own 10-K, 10-Q, 6-K, 424B4 or release, except SK hynix (its 10.28 release, read through a search copy) and Broadcom's 2025 AI figure (from its earnings call, reported); Oracle's columns are fiscal years ending in May; world chip sales from SIA and WSTS, read through search copies of their releases, 2026 a forecast. Backlog already signed but not yet revenue: Oracle $664 billion; Microsoft $677 billion; CoreWeave $132.5 billion. A promise of future revenue, not cash. VERIFIED

AI WENT FROM A SIDE BUSINESS TO MOST OF THE BUSINESS

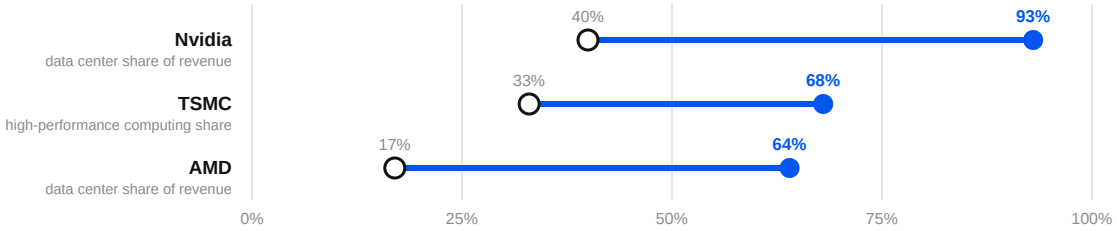

Open circle: 2020. Blue: latest quarter, 2026. Company filings; Nvidia FY2021 against Q3 FY2027; AMD 2020 as recast.

The money loop

Money moves down the stack: from the people and companies who pay for AI, to the labs, to the clouds that rent out the machines, to the chip designers, to the factories that make the chips and the memory. A growing share now moves back up as credit and equity, as the suppliers lend to, guarantee and invest in their own customers. ANALYSIS

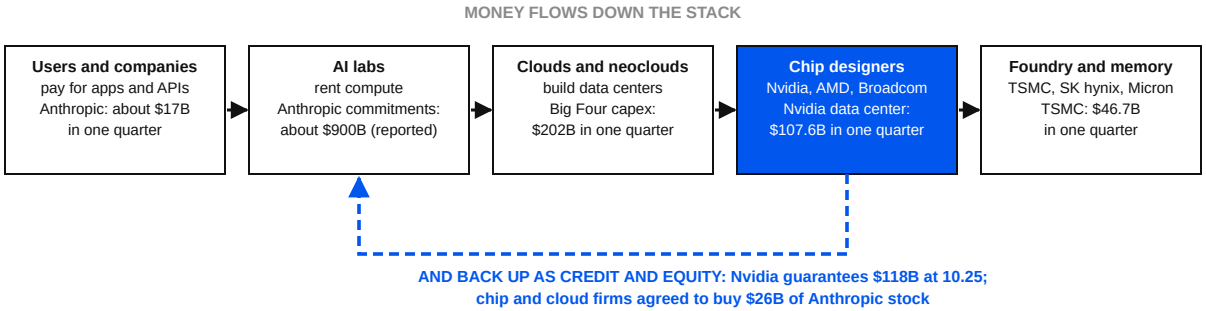

Latest quarter figures from company filings and releases; Anthropic commitments as reported; Nvidia guarantees from its CFO commentary.

Nvidia's 10-Q for the quarter to July showed $108.5 billion of guarantees, $105 billion of them for an Ohio campus leased to an OpenAI affiliate; its CFO commentary for the quarter to 10.25 puts maximum exposure at $118.0 billion, with about $32 billion of neocloud deals that hold only if partners find financing. Nvidia, Google, AMD, Broadcom, Micron, Microsoft and TSMC agreed to buy $26.0 billion of Anthropic stock at $83 alongside its IPO, per the prospectus. VERIFIED **Takeaway:** the money runs down the stack, and a growing share runs back up as credit.

AI and the economy

The build-out is no longer a corner of the tech sector. On Apollo Global Management's figures, the hyperscalers' capital spending rose from 0.3% of U.S. GDP in 2019 to 1.4% in 2025, already above the telecom peak of about 1.2% in 2000, and consensus estimates put it near 3% a year from 2027. CLAIM

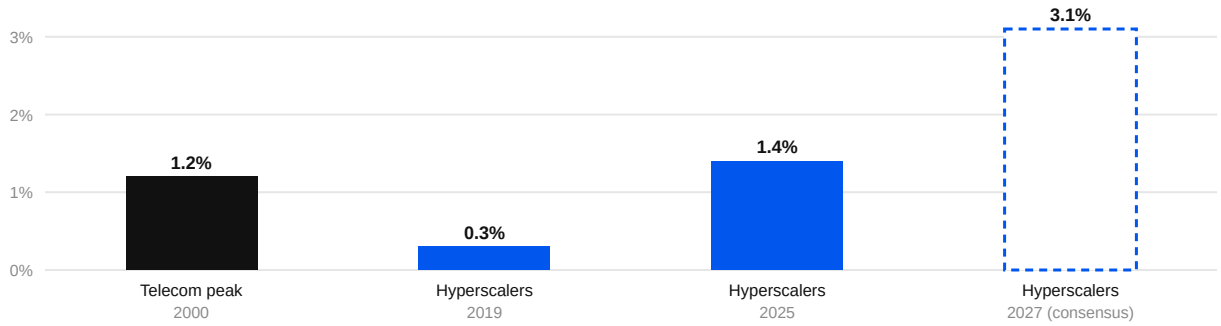

Capital spending as a share of U.S. GDP. Apollo Global Management (Torsten Slok), from consensus estimates; 2027 is a forecast, dashed.

It is also carrying a large share of growth. Federal Reserve staff estimated, in a note of 07.17, that spending on software, data centers, power facilities and computers added about 0.73 percentage point to annualized growth in the first quarter of 2026, after netting out computer imports; growth was then estimated at 2.1% and has since been revised to 2.5%. Real GDP grew at a 3.4% rate in the third quarter, with investment among the drivers, per the Bureau of Economic Analysis's advance estimate. VERIFIED



Fed staff note, 07.17 (FEDS Notes), net of computer imports; growth as then estimated, since revised to 2.5%.

Two links run from here to the wider economy, and both point back to financing. If the spending slows, a visible slice of U.S. growth slows with it; estimates of how big that slice is vary widely with the method, and the official accounts have no single "AI" line. And the newer builders, the neoclouds and data center developers, borrow at rates that move with credit markets: CoreWeave's 2032 notes yielded above 16% on 11.06 (reported). The telecom boom ended when the builders' financing broke, not when the traffic stopped. ANALYSIS

The next readings: the second estimate of third-quarter GDP on 11.25 (BEA) and the Federal Reserve's meeting on 12.08 to 12.09. Context, not cause: these charts show size, not a forecast. ANALYSIS

Two ways a chip boom ends

The industry trips on itself. These are the micro busts: too much inventory, falling prices. They come fast and hit memory hardest. Nvidia's revenue fell 24% in the quarter to January 2019, with gaming down 45%, as sales slowed after the crypto-mining boom; in fiscal 2023 it took $2.17 billion in inventory charges. Micron's revenue fell 49.5% in fiscal 2023. VERIFIED Down phases like these have averaged about 12 to 15 months (reported). CLAIM

The money or the economy turns. These are the macro busts, and they go deeper. World chip sales fell about 32% in 2001 and about 9% in 2009. The 2001 case is the closest match to today: a few companies built networks ahead of their own revenue, largely with borrowed money. North American carrier spending fell from about $97 billion in 2000 to about half that by 2002, by RHK's estimates (reported). Cisco took a $2.25 billion inventory charge in April 2001. Internet traffic kept growing through the bust; what broke was the financing of the builders. CLAIM

WARNING SIGN	WEIGHT	SEEN BEFORE	READING NOW	BEFORE 11.12	AS OF 11.22
Buyers slow their spending	High	2001; 2023 at Amazon and Meta	All four raised spending 80% to 120% from a year earlier	● Clear	● Clear
Spending outruns the buyers' own revenue	High	Telecom, 2000	$202 billion of Big Four capex in a quarter, well above the cloud revenue three of them report	◖ Watch	◖ Watch
Suppliers lend to and guarantee their customers	High	Telecom vendors, 2000 to 2001	Nvidia guarantees $108.5 billion at July, $118.0 billion at 10.25; chip and cloud companies agreed to buy $26 billion of Anthropic stock	◖ Rising	◖ Rising
Lenders demand more	High	2001, 2008	CoreWeave's 2032 notes yielded above 16% on 11.06 (reported)	▲ Flashing	▲ Flashing; no new reading
Prices and margins fall	Medium	Memory, 2019 and 2023	Nvidia guides gross margin to 71% to 72%, from 74.8%	● Clear	◖ Watch
Order books stop turning into revenue	Medium	Telecom, 2001	Backlogs still growing; no reported cancellations	● Clear	● Clear

Status is our judgment from the filings and reports cited. ● Clear, ◖ Watch or Rising, ▲ Flashing. ANALYSIS

IS DEMAND AT ITS PEAK?

On the record so far, no: the buyers are still raising spending and the suppliers are still sold out or growing. The harder question, whether it can last, is answered in the financing. If the next downturn comes, it is more likely to start with the lenders than with the buyers: the 2001 pattern, not the 2023 one. ANALYSIS

Does the industry depend on one man?

Less in orders than in financing and exports, and more than its filings say. Nvidia's annual report says it is "highly dependent on the services of our longstanding executive team," without naming Huang. VERIFIED

AREA	DEPENDS ON ITS FOUNDER	EVIDENCE
This year's orders	Low	Order books are contracts; the guide still rose. No changed orders reported.
The roadmap	Low	Vera Rubin began shipping on the existing plan, per the company.
Supply of chips and memory	Low	Multi-year supply commitments of $390 billion, up from $279 billion, per its CFO commentary.
Financing its customers	Medium to high	The guarantees and contingent neocloud deals sit with an interim chief executive who is also the finance chief.
Governments and exports	High	Nvidia's guide assumes no data center compute sales to China, though $2.4 billion of Hopper products shipped there last quarter, per its CFO commentary; Commerce's remote-access rule published 11.19, effective that day.
The bet after Rubin	Medium	Architecture and pricing calls are made a year or more out; no permanent chief executive is named.

Huang held 3.58% of Nvidia's shares as of 03.23.2026, per the proxy, including foundation shares in which he has no financial interest. No longer an officer or director, he may file fewer ownership reports. VERIFIED

Deals then and next

The founder era. Nvidia sold chips, then helped its customers pay for them: equity stakes in labs and neoclouds, capacity it agreed to buy back, and guarantees of its customers' leases. ANALYSIS **What may come next.** An interim

chief executive who is also the finance chief may weigh the balance sheet more heavily than new guarantees, and TSMC is setting credit limits for AI customers, Taiwan's Commercial Times reported (TSMC has not confirmed). Deals may shift toward cash up front and fewer contingent commitments. SPECULATION

From prompt to chip: what happens when you press Enter

Every chatbot answer makes the same trip: down through software, into a building full of chips, and back up as words, one word piece at a time. The labs differ mainly in which layers they own and which they rent. ANALYSIS

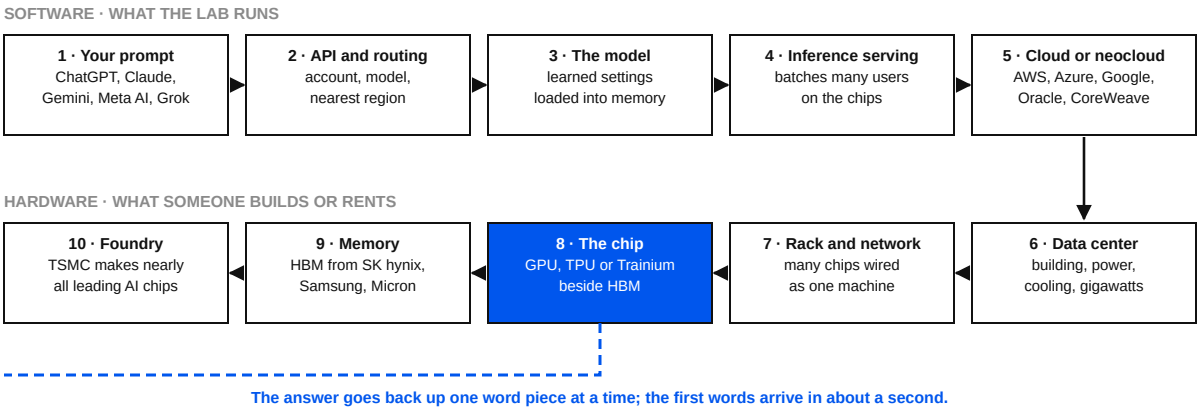

Built months ahead: the foundry makes the chip long before your prompt. Every other layer runs live.

LAYER	OPENAI	ANTHROPIC	GOOGLE	META	XAI	MICROSOFT
App, model, serving	Own	Own	Own	Own	Own	Own models; hosts OpenAI's
Cloud	Azure, Oracle, CoreWeave, AWS	AWS, Google Cloud, Azure	Google Cloud	Mostly own (reported)	Own	Azure
Data center	Stargate sites with partners	Partners' sites, plus its own U.S. sites with Fluidstack	Own	Own campuses	Colossus, Memphis	Own, plus leases
Chip design	Nvidia, AMD; own chip with Broadcom	Trainium, TPU, Nvidia	TPU, with Broadcom	Mostly Nvidia; own MTIA	Nvidia	Nvidia; own Maia
Manufacturing	TSMC	TSMC	TSMC	TSMC	TSMC	TSMC

From company announcements and press reports; where a lab has not disclosed its setup, treat these as reported. xAI and Microsoft added for completeness. CLAIM

WHERE THE SUPPLY NARROWS

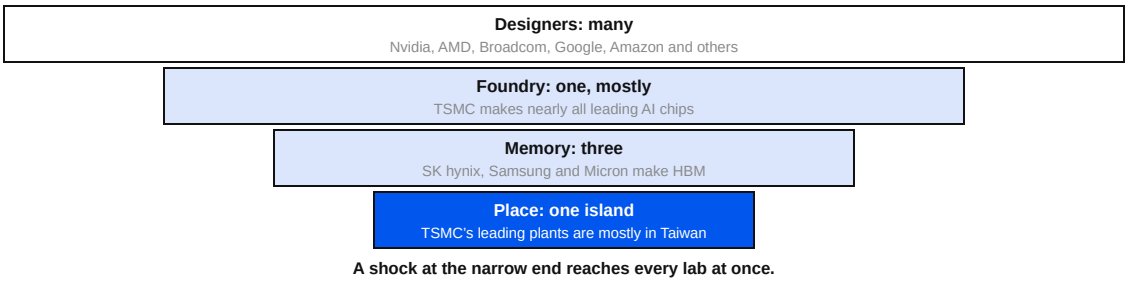

A labeled diagram, not a chart: verified shares for every layer were not available.

Words you need

GPU and accelerator	Chips that do many simple calculations at once. GPUs are Nvidia's and AMD's; custom accelerators such as Google's TPU and Amazon's Trainium are designed by a cloud company for its own data centers and rented to outside customers.
HBM	High-bandwidth memory: stacks of memory chips beside the processor in one package, made by SK hynix, Samsung and Micron.
Foundry	A factory that makes chips others design. Nvidia, AMD and Broadcom design; TSMC makes.
Hyperscaler	The giant clouds: Amazon, Microsoft, Google, and increasingly Oracle.
Neocloud	A newer cloud built to rent out AI chips, often with borrowed money: CoreWeave, Nebius, Lambda.
Training and inference	Training teaches a model once, at great cost. Inference runs it every time someone asks a question.
Data center and gigawatt	The building that holds the chips. AI capacity is now counted by the power it draws; a gigawatt is roughly a large power plant.
Capital spending	Money spent on buildings and equipment, as opposed to day-to-day costs. The single best gauge of AI chip demand.
Backlog	Work customers have signed for that the company has not yet counted as revenue. A promise of future revenue, not cash.

RED TEAM

The first crack shows in credit, not orders: by mid-2027 a major neocloud or data center borrower restructures its debt or loses financing that a chipmaker backed. Supporting it: neocloud bond yields above 16%, about $32 billion of Nvidia neocloud deals that depend on outside financing, and buyers' spending well ahead of their cloud revenue. **Wrong if** CoreWeave's yields fall back below 10% and the contingent deals close. ANALYSIS

BLUE TEAM

Nvidia's data center revenue grows at least 50% from a year earlier in each of the next two quarters, whoever is chief executive. Supporting it: the $130 billion guide for the January quarter, all four big buyers raising spending 80% or more, and record sales at SK hynix and Micron. **Wrong if** a big buyer cuts its 2027 spending plan in late-January results, or Nvidia guides below that pace in February. ANALYSIS

WHAT TO WATCH

11.25	Second estimate of third-quarter GDP, with investment detail (BEA)
12.01	Marvell results
12.08 to 12.09	Federal Reserve meeting: the cost of the money behind the build
12.09	Broadcom results
12.17 (expected)	Micron results
Late January	Big Four results and 2027 spending plans
Late February (date not yet announced)	Nvidia's fiscal fourth quarter, the first full quarter under interim leadership
No date	A permanent Nvidia chief executive

Dates as companies and agencies have announced them or as reported; the BEA and Federal Reserve dates were checked at their sources on 11.22, and the company dates are as the companies or reports announced them. Nvidia 8-K and Q3 FY2027 release and CFO commentary, SEC EDGAR Reg; Meta, Oracle, Broadcom, Micron, SK hynix and TSMC releases and filings Reg; Microsoft, Alphabet, Amazon, AMD and CoreWeave 10-Qs, Anthropic's 424B4 Reg; SK hynix release (search copy) and Broadcom's earnings call, reported Press; BEA GDP advance estimate, Q3 2026 Gov; Federal Reserve FEDS Notes, 07.17.2026 Gov; Apollo Global Management on capex and GDP Press; SIA historical sales

and WSTS forecasts, read through search copies Data; Nvidia 10-Q for the quarter to January 2019 and FY2023 10-K, Micron FY2023 10-K, Cisco 10-Q for the quarter ended 04.28.2001 Reg. How this desk is sourced: 16 labeled statements of fact, 10 verified against the cited primary document and 6 named reports labeled Claim; every company figure in the stack table was checked against the company's own filing or release on 11.21.2026, except as marked. A further 12 Analysis and 1 Speculation labels mark our own judgments. Counted label by label.

SAME QUARTERS, AND ONE DEAL READ TWO WAYS

Thursday's comparison of the two largest listings set Anthropic's September quarter against SpaceX's March quarter, the latest in SpaceX's prospectus. SpaceX has since filed two quarterly reports. Here are the same quarters side by side, and why the two companies' filings describe their compute deal at such different sizes.

Sources: SpaceX's 10-Q reports for the June and September quarters and their earnings releases, both companies' final prospectuses, and named press reports, labeled Claim. Follows section 7 of Special Coverage in Edition 051 (11.19). Mind Club research is developed through a human-led AI workflow using specialized expert agents, cross-model analytical challenge, primary-source review and substantial editorial direction. Independent analytical perspectives are used to test assumptions, identify weaknesses and refine conclusions. The human editor retains responsibility for the final work product. Some AI companies covered here supply models used in that workflow. Mind Club takes no position on any company or security.

THE SHORT VERSION

On the same quarters, the gap is much smaller than Thursday's comparison suggested, and the profit picture flips. Anthropic's September-quarter revenue was $17.0 billion against SpaceX's $14.6 billion, not more than three times as much. SpaceX made an operating profit of $2.8 billion that quarter; Anthropic made an operating loss of $567 million. The compute deal is not a contradiction: SpaceX described one agreement signed in May, and Anthropic describes that agreement plus amendments and new capacity through October. ANALYSIS

1 · The same quarters, side by side

SpaceX's revenue nearly doubled between the June and September quarters, and most of the increase came from its AI unit, which it now reports as Super Intelligence (SI). In the September quarter that unit's revenue was $7.6 billion, more than its Starlink business, and $7.0 billion of it came from its SI solutions and infrastructure line, mostly cloud computing services but including $608 million from Cursor. One unnamed customer, Customer B, accounted for 34.5% of SpaceX's total revenue that quarter; SpaceX does not say who it is. VERIFIED

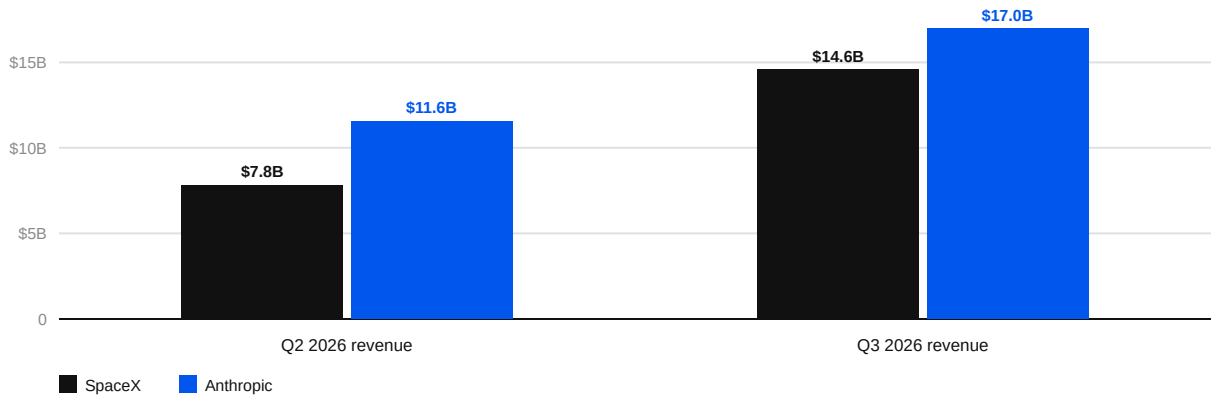

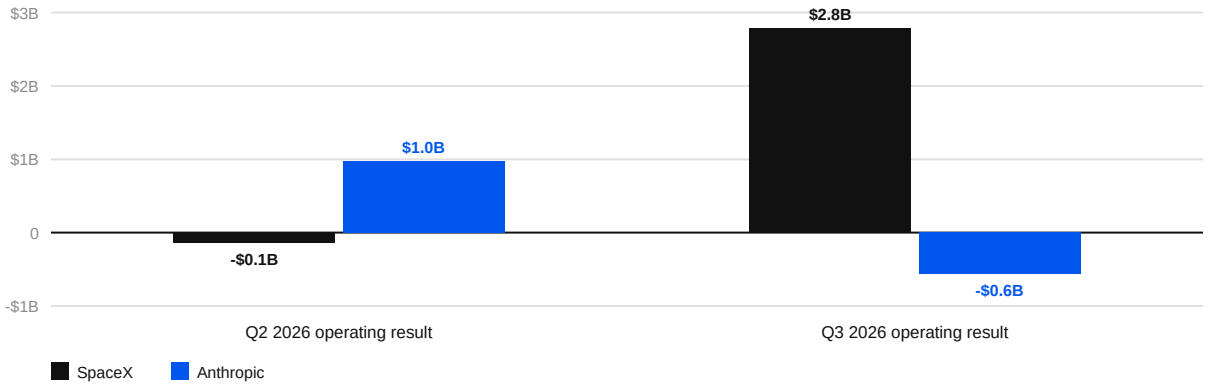

Quarterly revenue and operating income (loss), $ billions. SpaceX 10-Qs for the quarters ended 06.30 and 09.30.2026; Anthropic final prospectus, quarterly results. VERIFIED

$ MILLIONS	SPACEX Q2	SPACEX Q3	ANTHROPIC Q2	ANTHROPIC Q3
Revenue	7,814	14,575	11,555	16,956
Of which AI (SpaceX's AI, later SI, segment)	2,561	7,560	all	all
Of which SI (earlier AI) solutions and infrastructure line, mainly compute	2,194	7,022	n/a	n/a
Operating income (loss)	(143)	2,791	974	(567)
Net income (loss)	(541)	4,404	(66,309)	(23,291)
Capital spending (purchases of property and equipment)	18,369	19,593	925 for nine months, plus 1,955 of prepayments for leased assets	
Largest single customer, share of revenue	Customer B, 19.5%	Customer B, 34.5%	none above 10%; top five about 25% of nine-month revenue	

Anthropic's net losses are driven mostly by changes in the fair value of convertible notes, derivative liabilities and SAFEs ($59.6 billion in Q2 and $21.9 billion in Q3), a non-cash charge. SpaceX's Q3 AI infrastructure revenue includes $608 million from Cursor. Anthropic does not report capital spending by quarter. VERIFIED

What changes from Thursday. Edition 051 said Anthropic's September-quarter revenue was more than three times SpaceX's March-quarter revenue, the latest checked then, and flagged that the periods differed. On the same quarter the ratio is about 1.2 times. SpaceX's capital spending, about $19.6 billion in the September quarter alone, is more than twenty times what Anthropic spent on property and equipment in nine months, and about seven times that figure with its $1.955 billion of lease prepayments added; Anthropic's spending runs through leases and commitments instead. ANALYSIS

2 · One deal, two descriptions

SpaceX's prospectus, filed 06.12, says Anthropic agreed in May "to pay us $1.25 billion per month through May 2029," with capacity ramping in May and June at a reduced fee, terminable by either side on 90 days' notice after an initial three months. Anthropic's prospectus, filed 11.19, says that between May and October it "entered into and amended several agreements" with SpaceXAI, agreeing to spend "up to $120.5 billion" over terms extending into 2029, about $3.1 billion a month once all capacity is delivered, expected in December, with non-cancelable periods ending in January and March 2027. VERIFIED

TERM	SPACEX'S DESCRIPTION	ANTHROPIC'S DESCRIPTION
Filed	Prospectus, 06.12.2026	Prospectus, 11.19.2026
What it covers	Agreements signed in May 2026 (one agreement dated 05.03 in the financial statement notes)	Several agreements entered into and amended between May and October 2026
Capacity	About 325,000 Nvidia GPUs across COLOSSUS and COLOSSUS II	Compute powered by Nvidia GPUs; full delivery expected December 2026
Price	$1.25 billion a month, reduced fee during the ramp	About $3.1 billion a month at full delivery; different fees during ramps
Term and total	Through May 2029; no total stated	Into 2029; up to $120.5 billion in total
Exit	90 days' notice after an initial three months	90 days' notice after non-cancelable periods ending January and March 2027

What the arithmetic says. $1.25 billion a month for the 36 months from June 2026 to May 2029 comes to about $45 billion, which matches the "nearly $45 billion" Bloomberg reported from SpaceX's filing in May. ANALYSIS At full delivery the expanded deal runs about $1.85 billion a month above the original one. At $3.1 billion a month, the 30 months from December 2026 to May 2029 come to about $93 billion. Undisclosed ramp fees before December would have to cover the rest of the $120.5 billion, or some terms run past May 2029; neither filing says which. ANALYSIS

What the filings support. Read together, the two documents describe a deal that grew: a May agreement, then amendments and added capacity through October, with longer lock-in periods. SpaceX's own September-quarter report says it signed "additional contracts" for cloud services in the quarter. The two numbers describe different stages of the same relationship, not a disagreement. ANALYSIS A press report in September put the total at up to $84.5 billion, citing confidential IPO documents relayed in a social media post; we could not check it. CLAIM It may reflect an earlier stage, before the October amendments. ANALYSIS

3 · What stays open

QUESTION	WHAT WE KNOW	WHAT WOULD SETTLE IT
Is Customer B Anthropic?	SpaceX does not name it. It is the only customer above 10% whose revenue all falls in the AI segment, and its share jumped from 19.5% to 34.5% as the deal ramped. Probable, not disclosed. ANALYSIS	SpaceX naming the customer, or Anthropic's first 10-Q
How much of the $120.5B is firm?	Either side can leave on 90 days' notice after January and March 2027. SpaceX reported $14.1 billion of contracted sales from several cloud agreements in its June-quarter release, a figure it defines as the non-cancellable, enforceable period; that is far below the headline totals. ANALYSIS	SpaceX's year-end backlog disclosure
When does each agreement end?	SpaceX: through May 2029. Anthropic: "into 2029," on respective schedules.	Either company filing the agreements or their key terms

QUESTION	WHAT WE KNOW	WHAT WOULD SETTLE IT
Anything new since 11.19?	No change since 11.19: no new SpaceX or Anthropic filing on EDGAR through 11.22, and no statement found	A 10-Q, 8-K or amended prospectus from either

WHY IT MATTERS

The deal ties two newly public companies together. A large share of SpaceX's fastest-growing revenue line appears to rest on one customer that can leave on 90 days' notice from early 2027, and a large share of Anthropic's computing power rests on a supplier that also builds a rival model, Grok. Each now depends heavily on the other. ANALYSIS

RED TEAM

A quarter is a short window. SpaceX's September quarter was the first full quarter after its cloud agreements ramped in May and June, and Anthropic's operating result has swung from loss to profit and back to loss over three quarters, so neither quarter may represent the year. Treating Customer B as Anthropic is an inference, and the deal arithmetic rests on fees neither company fully discloses. ANALYSIS

SpaceX 10-Q, quarter ended 06.30.2026, filed 08.04 REG; SpaceX 10-Q, quarter ended 09.30.2026, filed 10.22 REG; SpaceX earnings releases (8-K Ex. 99.1), 08.04 and 10.20 REG; SpaceX final prospectus (424B4), 06.12 REG; Anthropic final prospectus (424B4), 11.19 REG; Bloomberg Government, 05.20 PRESS; CoinGape, 09.30 PRESS. How this desk is sourced: 5 labeled statements of fact, 4 verified against the cited primary document and 1 named report labeled Claim; a further 10 Analysis labels mark our own judgments. Counted label by label.

04 / AI THEME PULSE — NINE THEMES · EVIDENCE VS COVERAGE · 11.16 TO 11.22

AI THEME PULSE

Scientific AI and frontier models ran warm on the evidence. Chips split hardest: Nvidia's results read well, the coverage of its week did not. Capital split too, with record raises on the record and a sell-off in the news.

ANALYSIS

THEME	EVIDENCE	COVERAGE	ITEMS	CONFIDENCE	WHAT DROVE IT
Frontier models	+1.0	0.0	4 · 1	LOW	Four releases with system cards or model pages; outside scores split. SPLIT
Agents & automation	QUIET	QUIET	1 · 0	·	Fewer than three scored items this week
Compute & infrastructure	0.0	−1.0	2 · 1	LOW	The BIS remote-access rule and reports of peaking capex growth outweigh Blackstone's data center buy. SPLIT
Chips & memory	0.0	−1.7	2 · 3	LOW	Nvidia's $130 billion guide against a reported SEC inquiry, an unfiled 10-Q and a 22.92% weekly fall. SPLIT
Enterprise adoption & pricing	QUIET	QUIET	2 · 0	·	Fewer than three scored items this week
Safety, security & incidents	−0.3	−1.0	3 · 1	LOW	OpenRouter's provider removals and the GPT-6.2 card's blocker finding outweigh OpenAI's distillation defenses.

THEME	EVIDENCE	COVERAGE	ITEMS	CONFIDENCE	WHAT DROVE IT
Policy, regulation & courts	−0.5	−1.0	2 · 1	LOW	A rule in force, a rescinded rule and China's warning; the FDA's wording change scored neutral.
Scientific AI & AI x biotech	+1.0	+1.0	4 · 1	LOW	A Lean-checked Riemann proof, a peer-reviewed RNA design agent and a Merck deal; MedChemBench scored neutral.
Capital & deals	+1.0	−1.0	3 · 2	LOW	A record IPO and placement and two big rounds, against a falling ANT and Bloomberg's wipeout framing. SPLIT

−2 NEGATIVE +2 POSITIVE QUIET: FEWER THAN 3 ITEMS

Not scored: Dragos's provider report (11.19) and NebuSec's Muse write-up (11.20), both incident disclosures in the window; scoring them would likely lower the Safety row and may lift Agents & automation out of QUIET.

Method: every dated event from a primary source (releases with their Claim Check verdict, pricing changes, incident disclosures, rules, rulings and filings, financings of $100M or more and M&A) and every open news or analysis article is scored −2 to +2 for its stance on the theme's progress, never on a share price. Scores are averages; Items counts events · articles. SPLIT marks rows where evidence and coverage disagree in sign or by 1.0 or more; ONE OUTLET marks rows where one publisher supplies most of the coverage. ANALYSIS

05 / COMPUTE & CAPITAL MAP FIVE WEEKS OF AI FINANCINGS ≥$100M AND M&A · 2026 YTD

COMPUTE & CAPITAL MAP

Eleven deals for $37.3B, five of them verified. One line, Anthropic's $26.0B placement with chip and cloud companies, is about 70% of the week; without it, the week came to about $11.3B. ANALYSIS

WEEK	DEALS	TOTAL	VERIFIED	LARGEST
10.19 to 10.25	13	$10.5B	6	Salesforce and Braze $5.66B · Tenstorrent $1.5B · Flapping Airplanes $1.0B
10.26 to 11.01	8	$2.4B	3	Bedrock Robotics $600M · Itaú Unibanco and NeoSpace $589M · Upwind Security $400M
11.02 to 11.08	13	$137.0B	2	Marvell and Lumentum $115.0B · OpenAI $11.0B · Shield AI and Echodyne $5.3B
11.09 to 11.15	10	$5.3B	7	Mirendil $1.0B · Mercor $750M · Fractile $720M
11.16 to 11.22	**11**	**$37.3B**	**5**	**Anthropic Class N placement $26.0B · Isomorphic Labs $5.0B · Palo Alto Networks and Applied Compute $4.5B**

THIS WEEK BY SECTOR	TOTAL	DEALS (CHANNEL, STATUS)
Foundation models	$26.0B	Anthropic Class N placement $26.0B (Strategic, C)
Healthcare & scientific AI	$5.31B	Isomorphic Labs $5.0B (Growth, V) · MegaRobo about $200M (Growth, C) · Boltz $110M (Venture, V)
Agents & dev tools	$4.61B	Palo Alto Networks and Applied Compute $4.5B (M&A, V) · Noeri $110M (Venture, C)
Enterprise & vertical apps	$706M	Payhawk $280M (Growth, V) · Model ML $250M (Venture, V) · Blue J $176M (Growth, C)
Robotics & autonomy	$500M	Genesis AI $500M (Venture, C)
Cybersecurity	$125M	Corma $125M (Venture, C)

By channel this week: Strategic $26.0B · Growth $5.66B · M&A $4.5B · Venture $1.10B. Prior four weeks averaged $38.8B, lifted by Marvell's deal of about $115B. Not counted: Anthropic's $81.34B public offering (not a tally channel), reported talks at Cognition, Databricks and Rogo, Nvidia's reported deal for Inception (not announced), and funds such as Norwest's $4B and Plural's €1B. The placement is marked C: it closed with the IPO per Edition 052, and we did not reopen the closing filing. Capex guidance is never added.

2026 YEAR TO DATE, BY SECTOR AND CHANNEL

SECTOR	VENTURE	GROWTH	STRATEGIC	M&A	INFRA CAPEX (GUIDE)	TOTAL TRACKED (EX-CAPEX)
Foundation models	$23.8B	$251.5B	$31.0B	$8.2B	·	$314.5B
Chips & semis	$4.3B	$1.8B	·	$119.0B	·	$125.0B
Agents & dev tools	$2.0B	$2.5B	·	$64.5B	·	$69.1B
Enterprise & vertical apps	$1.2B	$5.7B	·	$38.7B	·	$45.6B
Infrastructure & compute	$3.2B	$14.2B	·	$24.0B	$725B+	$41.4B
Robotics & autonomy	$3.6B	$24.2B	·	$10.7B	·	$38.6B
Cybersecurity	$1.0B	$950M	·	$32.5B	·	$34.5B
Healthcare & scientific AI	$2.5B	$8.1B	·	$1.2B	·	$11.8B
Data infrastructure	$300M	$11.2B	·	·	·	$11.6B
Consumer AI	$2.5B	·	·	·	·	$2.5B
TOTAL TRACKED	$44.4B	$320.3B	$31.0B	$298.8B	$725B+	$694.5B

LESS MORE 155 transactions (94 verified). Square-root shade scale.

VC watch. Plural announced a €1B fund for European frontier technology on 11.16; Greenoaks and Sequoia led Isomorphic's round, General Atlantic and QED Model ML's, and Bessemer Boltz's. VERIFIED Cognition and Databricks are reportedly raising at higher prices despite the sell-off. CLAIM

RED TEAM

A placement bought by the IPO company's own suppliers is not the same kind of money as a venture round. Counted as strategic, it makes this week look far bigger than the private market was.

06 / BENCHMARK REALITY CHECK WHO MEASURED · WHICH TEST · WHAT THE DOCUMENT SAYS

BENCHMARK REALITY CHECK

Nine headline numbers from the week's releases. Seven were measured by the people who made the thing; one outside check is a proof checker, not a benchmark, and one is an independent index score we could not find on the index's own site.

CLAIM	WHO MEASURED	TEST	VERDICT
The Riemann proof's statement matches Mathlib's and several independent checkers accept the proof	Lean community, MIT	Lean kernels and independent checkers	MATCH
GPT-6.2 Astra scores 5.5 points above Claude Fable 5.2 on Agents' Last Exam at about 5% of the cost	OpenAI	Agents' Last Exam; nine Fable tasks excluded	PARTIAL
GPT-6.2 Astra gains 9 points on Terminal-Bench-Science over GPT-6 Astra	OpenAI	Terminal-Bench-Science 0.1; rival on another harness	PARTIAL
GPT-6.2 Astra gains 6.2 points on GDPval v2 with 22% fewer output tokens	OpenAI	Its own internal test	OPEN
Gemini 4 Neon scores 61 on the Artificial Analysis Intelligence Index, third overall	Google, quoting the index	AA Intelligence Index, at its highest thinking setting	OPEN
Gemini 4 Neon scores 39.0% on SWE-sweep	Google	Its own harness; 12.8% on a neutral one	PARTIAL
MAI Code 2 delivers 96% of Claude Opus 5.5's performance at 5% of its cost	Microsoft	Internal MSBench, promotional prices	OPEN
Reasoning extraction falls from 54.0% to 0.2% with new safeguards	OpenAI	Internal red-teaming	OPEN
GPT-6.2 Astra scores 84.5% on MedChemBench	OpenAI	Its own new benchmark	PARTIAL

Partial: the number is in the document, but the comparison has a stated catch: excluded tasks, mismatched harnesses, a harness that triples the score, or possible exposure to the source papers. Open: the test is internal, gives no absolute score, or we didn't find the number on the index's own site. The GPT-6.2 system card also says both models were more likely than GPT-6 Astra to leave out obstacles that blocked their work.

AI X BIOTECH REALITY CHECK

CLAIM THIS WEEK	REGISTRY	WHAT THE REGISTRY SHOWS, 11.22	VERDICT
Isomorphic Labs is "setting the stage for clinical entry" after a $5B raise	ClinicalTrials.gov sponsor search	No trial registered under Isomorphic Labs	OPEN
Recursion has FDA alignment on a registrational Phase 3 for REC-4881 in FAP	NCT05552755	Phase 1b/2 TUPELO recruiting, enrollment 67; the Phase 3, planned for the second half of 2027, is not yet registered	PARTIAL
Insilico's rentosertib is in a Phase 3 of 320 patients at 47 centers	NCT07687459	Phase 3, recruiting, 320 planned, 47 locations	MATCH
Iambic's IAM1363 reached the clinic	NCT06253871	Phase 1/2, recruiting, 463 planned	MATCH

Reality roundup. Four claims checked against the registry: two match, one is partly ahead of the record, and one has nothing to check yet. A registered trial says a drug is being tested, not that it works; "world's first" labels are the companies' own. REG DATA

TRY THIS WEEKEND

Three small, safe tests on public or dummy material. Scan the last column for the ones built for biotech and pharma teams.

#	EXPERIMENT	TIME	DATA YOU NEED	BIOTECH & PHARMA
A	**Price one real task, not a benchmark**	30 min	One non-confidential task you do every week	General
B	**Find out who actually runs your model**	15 min	Your own account settings	General
C	**Check an AI drug claim against the registry**	20 min	A press release and ClinicalTrials.gov	BIOTECH & PHARMA

A Price one real task, not a benchmark

TRY IT WITH	Two models you can already use, such as GPT-6.2 Sol and Gemini 4 Neon, each at its default model settings, on an account where you have checked whether your chats are kept or used for training
WHAT IT TESTS	Whether a vendor's cost-per-task claim holds for your own work
SETUP	Pick one multistep, non-confidential task. Give both models the same prompt and note the answer, the time and the tokens or cost each one shows.
YOU'LL KNOW IT WORKED WHEN	You can say which answer you would use and what it cost, in your own numbers
BIOTECH & PHARMA VERSION	Summarize a public trial result from its registry page and press release, then check every number against the source.
GUARDRAIL	Public material only. No patient data, no unpublished results, nothing under confidentiality.

B Find out who actually runs your model

TRY IT WITH	Your account on any model router or cloud platform you already use
WHAT IT TESTS	Whether your prompts can go to providers you never picked, and under what data terms
SETUP	Open the provider and data settings. List which providers can serve your requests and whether each is marked zero-retention. Make no changes until you have read the list.
YOU'LL KNOW IT WORKED WHEN	You can name every provider your settings allow, and you have chosen, on purpose, which ones may. A zero-retention mark is the provider's promise, not proof.
BIOTECH & PHARMA VERSION	Do this with your IT or security team before any regulated or confidential data goes through a router.
GUARDRAIL	Settings changes can break running work. Change one thing at a time, note the old value so you can switch back, and on a work account let your IT team decide. Never paste a real key or credential into a test prompt.

C Check an AI drug claim against the registry

TRY IT WITH	Any assistant that can read a web page, and ClinicalTrials.gov
WHAT IT TESTS	Whether an assistant reads a registry record correctly, and whether a company's claim matches it
SETUP	Take a recent release about an "AI-discovered" drug. Ask the assistant for the trial number, phase, status and enrollment, then open the record yourself and compare.
YOU'LL KNOW IT WORKED WHEN	Every field matches the record, or you have found where the release runs ahead of it
BIOTECH & PHARMA VERSION	This is the biotech version. Try it on the four rows in 06.

MIND CLUB IS PUBLISHED BY TORPEDO PUBLISHING LLC

08 / PAPER OF THE WEEK ONE PAPER · EVIDENCE LEVEL FIRST

PAPER OF THE WEEK

PEER "Programmable functional RNA design with a virtual scientist," Nature Biotechnology, published 11.19, from groups at the Chinese University of Hong Kong, MIT, the Wyss Institute and the Broad Institute, among others. PMID 43058249.

What they did. The authors built RhoAgent, an AI agent that designs RNA molecules with a diffusion model and refines them with feedback from lab experiments. VERIFIED

What they found. It produced working ribozymes, RNA molecules that act like enzymes, with only 40% to 70% sequence similarity to natural ones. It made one ribozyme cut up to six times more efficiently than an expert-trimmed version and improved a fluorescent RNA by 230%, according to the abstract. VERIFIED

Why it matters. Designing RNA that does a job, not just RNA that folds, is a step toward programmable RNA tools and medicines. A peer-reviewed agent that works with a lab loop is rarer than one that only works on a benchmark. ANALYSIS

RED TEAM

Lab assays on a handful of RNA systems, with no animal or therapeutic tests in the abstract. The gains depend on rounds of lab feedback, so the agent is only as fast as the bench behind it.

Nature Biotechnology, doi:10.1038/s41587-026-03362-3 PEER · PubMed 43058249 REG

09 / THE DECODER AI VOCABULARY, EXPLAINED SIMPLY · WEEK OF 11.16

THE DECODER

The week's fourteen words in one table, three older words for review, a five-question quiz and one deep cut. Words explained here are underlined like this the first time they appear in later pages.

WORD	IN ONE LINE	DAY
Patch diffing	Comparing a program before and after a security update to see what was fixed; the fix points attackers at the hole	11.16
Coordinated vulnerability disclosure	Telling a software maker about a bug privately and waiting an agreed time before going public, so a fix can ship first	11.16
Interim final rule	A government rule that takes effect before the public comments on it; comments can still change it	11.16
Gross margin	The share of each dollar of sales left after paying what it cost to make the product	11.17
Federal Register public inspection	A rule posted for anyone to read a day or more before it is officially published; most clocks start on publication	11.17
Agent identity	Giving an AI agent its own login, like a badge, so its actions can be logged, limited and switched off	11.17
Purchase commitment	A promise to buy a set amount from a supplier over years; it locks in supply but can leave the buyer paying for what it can't use	11.18

WORD	IN ONE LINE	DAY
Upsizing	Selling more shares than first planned, shortly before pricing, usually because demand is strong	11.18
CVSS score	Common Vulnerability Scoring System: a 0 to 10 score for how bad a security hole is; 9 and above is critical	11.18
Lock-up	A promise by pre-IPO owners not to sell for a set time; eligible to sell is not the same as sold	11.19
Opening auction	For a new listing, the exchange collects orders and finds the one price where the most shares trade; that is the first trade	11.19
Critical Capability Level	A line Google DeepMind draws for how dangerous a model's skill could be; near it, Google says it adds safeguards	11.19
Options expiration	The day a month's stock options run out, usually the third Friday; unwinding hedges can make prices swing more	11.20
Gross leverage	A fund's total bets, long and short added together, compared with its own money	11.20

SPACED REVIEW: THREE WORDS

SI (STANDING TERM)

SUPER INTELLIGENCE, US FEDERAL USAGE SINCE 09.29.2026

An executive order tells US federal agencies to call the AI tools you already use "Super Intelligence" in their own documents, where the law allows. Same tools, new name tag. Researchers still use "superintelligence" for a hypothetical AI that would beat people at almost everything.

BIOTECH BRIDGE A sponsor can call its program a "breakthrough" in its own deck, but only an FDA designation makes the word official. The molecule doesn't change either way.

SEEN IN THE BRIEF Three more agencies now use it, including the FDA (10).

INFERENCE PROVIDER

A company that runs AI models on its own computers for others. It may hold your keys, your data and sometimes your model.

BIOTECH BRIDGE A contract manufacturer: you own the recipe, they run the plant.

SEEN IN THE BRIEF OpenRouter removed three of them (01, item 03).

LEAN

A PROOF ASSISTANT (SOFTWARE)

A program that checks a mathematical proof step by step, so errors in the reasoning are very hard to slip through. It only checks the statement it was given.

BIOTECH BRIDGE A validated statistical analysis: it confirms the math, not whether the endpoint was the right one.

SEEN IN THE BRIEF The Riemann proof passed several Lean checkers (01, item 01).

DEEP CUT · FORM 12B-25

A public company that can't file its quarterly report, the 10-Q, on time can file a short notice, Form 12b-25, often called an NT 10-Q. It buys five more calendar days and must give a reason. Large accelerated and accelerated filers are due 40 days after the quarter ends (others 45), and the notice is due within one business day of the missed deadline. A late notice is not a finding of anything; it is a signal that the company isn't ready to sign. ANALYSIS

SUNDAY QUIZ Five questions, answers upside down below.

1. A rule takes effect today, and the public can comment on it for 60 days. What kind of rule is it?
2. A company sells more IPO shares than planned because demand is strong. What is that called?
3. Early owners can't sell their shares for 180 days after the listing. What is that promise called?
4. A fund holds $100 of stocks and is short $100 of others, on $100 of its own money. What is its gross leverage?
5. A proof checker accepts a proof. Does that mean mathematicians understand it?

1. An interim final rule. 2. Upsizing. 3. A lock-up. 4. 200%. 5. No: it means the logic holds for the statement given; explaining it is separate work.

THE A TO S ON INTELLIGENCE

The register's biggest week since it opened: the FDA, Commerce and the State Department all moved to the new terms on 11.15 and 11.16, taking the count from 15 to 18. Commerce's chip rule, published on 11.19, now carries "Super Intelligence (SI)" in the Federal Register itself.

SOURCE	TYPE	HOW IT'S USED	MANDATED OR VOLUNTARY	FIRST SEEN	BADGE
New 11.16: FDA, "FDA's Use of Super Intelligence and SI under Executive Order 14434": new public documents use the terms from 11.16; statutes and regulations keep "AI" as written; no new requirement for sponsors	Federal (agency)	Replaces AI	Mandated (cites the order)	11.16	GOV Verified
New 11.16: Commerce (BIS), remote-access rule and release: "advanced super intelligence (SI)" chips and servers; the rule as published 11.19 (91 FR 74970) says in footnote 1 that it uses "Super Intelligence (SI)" in place of AI, citing EO 14434	Federal (agency)	Replaces AI	Mandated (cites the order)	11.16	GOV Verified
New 11.16: State Department, "The San Francisco Declaration" (11.15, nine partners): "super intelligence (SI)"; a footnote says that for some participants the term refers to "artificial intelligence"; the joint text does not cite the order	Federal (with foreign partners)	Replaces AI	Mandated, by our definition; the joint text doesn't cite the order	11.15	GOV Verified
New 11.22: Air Force Materiel Command, "Data and SI Strategy" implementation plan, per an AFLCMC news page we couldn't open; date unconfirmed	Federal (military)	Replaces AI	Mandated, if confirmed	Undated	GOV Claim
Defense Department, CDAO redesignated the "Super Intelligence Office" (SIO), per Breaking Defense; the department's release still wouldn't open for us on 11.22	Federal (executive department)	Replaces AI in an office name	Not determined (Claim)	11.13	PRESS Claim
NIH Notice NOT-OD-27-011, "USG Policy for the use of Super Intelligence": the government "is transitioning" to "Super Intelligence" and "SI" "pursuant to Section 2 of" EO 14434; keeps the statutory definition; names an NIH Office of the Chief Super Intelligence Officer	Federal	Replaces AI	Mandated (cites the order)	10.23	GOV Verified
the President's Truth Social post, "100% in favor of Super Intelligence (SI)," as AP reported (11.06)	Federal (presidential)	Replaces AI	The President's own usage, not counted	11.06	PRESS Claim
Justice Department, "Super Intelligence (SI) Task Force" (DOJ-SITF)	Federal (executive department)	Replaces AI in the task force's name	Mandated, by our definition; the release doesn't cite the order	11.05	GOV Verified
OCC Bulletin 2026-54, framework for examining banks' use of SI; SI not "a separate and distinct risk category"; non-binding examiner guidance that cites EO 14434 in a footnote	Federal (bank regulator)	Replaces AI in its own terms	Not determined	11.04	GOV Verified

SOURCE	TYPE	HOW IT'S USED	MANDATED OR VOLUNTARY	FIRST SEEN	BADGE
Tesla, chip roadmap renamed SI4 to SI6, per Tom's Hardware on a social post	Company	Replaces AI in product names	Voluntary	11.04	PRESS Claim
Commerce, "Super Intelligence Use Cases Inventory" page (wouldn't open; date unconfirmed)	Federal	Replaces AI	Mandated, if confirmed	11.08	GOV Claim
DOJ staff terminology memo, reported 10.06 (memo not public)	Federal	Replaces AI	Mandated, if confirmed	10.06	PRESS Claim
US Marine Corps message on SI terminology (not opened)	Federal	Replaces AI	Mandated, if confirmed	10.23	PRESS Claim
US Navy message (Navy Times, Stars and Stripes)	Federal	Replaces AI	Mandated, if confirmed	10.22	PRESS Claim
State Department cable, reported 10.20 (not public)	Federal	Replaces AI	Mandated, if confirmed	10.19	PRESS Claim
Tesla, 10-Q: "Automotive and SI Enabled Products"	Company (SEC)	Replaces AI in section and product names	Voluntary	10.21	REG Verified
SpaceX, Q3 results release: segment "Super Intelligence (SI)"	Company (SEC)	Replaces AI	Voluntary	10.20	REG Verified
IBM, Q3 results release: executive quotes	Company	Quoted only	Voluntary, not counted	10.21	CO Verified
SI.gov (AI.gov redirects there)	Federal	Replaces AI	Mandated	By 10.08	GOV Verified
HHS, "Super Intelligence at HHS"	Federal	Replaces AI	Mandated	10.09	GOV Verified
USDA, "Super Intelligence Compliance Plan"	Federal	Replaces AI	Mandated	Undated (checked 10.18)	GOV Verified
OSTP About page	Federal	Alongside AI	Mandated	Undated (checked 10.11)	GOV Verified
NIST, CAISI page headed "CAISSI"	Federal	Alongside AI	Mandated	Undated (checked 10.11)	GOV Verified
DOE Genesis Mission ("Super Intelligence" topic tag)	Federal	Alongside AI, tag only	Mandated	10.09	GOV Verified
FTC, ftc.gov/si "Super Intelligence Compliance Plan"	Federal (independent agency)	Replaces AI	Not determined	10.08	GOV Verified
Satya Nadella, "Models as Insider Risks" [Microsoft]	Executive, personal site	Replaces AI	Voluntary	10.10	CO Verified
Alpha Compute, 6-K exhibit self-description	Company (SEC)	Alongside AI	Voluntary	10.06	REG Verified

SOURCE	TYPE	HOW IT'S USED	MANDATED OR VOLUNTARY	FIRST SEEN	BADGE
NVIDIA, $1B science release	Company	Alongside AI	Voluntary	10.08	CO Verified
Micron, chief executive quote in its 8-K exhibit	Company (SEC)	Quoted only	Voluntary, not counted	09.30	REG Verified
Salesforce, "AIForce is officially SIForce" (chief executive's post)	Company	Replaces AI in a product name	Voluntary	10.09	Claim
"White House Accord on Super Intelligence" (six companies, per press)	Company	Not determined: no signed text found	Voluntary	09.29	PRESS Claim

Tally, 11.22: 18 sources verified (11 mandated: SI.gov, HHS, USDA, OSTP, NIST, DOE, DOJ, NIH, FDA, BIS, State; 5 voluntary: Nadella, Alpha Compute, NVIDIA, SpaceX, Tesla; 2 not determined: FTC, OCC). This week's changes: three verified rows added 11.16 (FDA, BIS, State), one Claim added (Air Force Materiel Command). Executive Order 14434 itself and the President's own usage aren't counted. No reversal found. Last verified change 11.16; Claim row added 11.22.

What would move it. The Defense Department's or Air Force's release opened at its source; a White House document carrying out the "ban" post (our call #21, due 11.30); "SI" in a company's risk factors; a non-US government using the term on its own. Under Sec. 3(b), proposed legislative language is due to the President by about 11.28; the order doesn't require it to be made public. Sec. 4(c) creates no enforceable rights, so "mandated" here means directed, not enforced. ANALYSIS

RED TEAM

The FDA says sponsors can still say AI and its rules keep the old words, so the switch is in what the agency writes, not what anyone else must do. More rows mean the label is spreading through government, not that the technology changed. We say AI.

Register rows as opened on their first-seen dates GOV; FDA page (11.16), BIS rule (91 FR 74970, 11.19) and State Department declaration (11.15) opened 11.16 and 11.17; Breaking Defense for the Pentagon office (11.13) PRESS; the AFLCMC page would not load 11.22.

11 / PATTERN RECOGNITION

TWO TO FOUR WEEKS · THREE PATTERNS

PATTERN RECOGNITION

PATTERN	THE EVIDENCE	STRENGTH
The chip trade is now a governance and credit trade	Nvidia's founder exit, a reported SEC inquiry, its guarantees to customers and neocloud bond yields, all moving shares more than the $130 billion guide	Rising
Chipmakers are funding their customers	AMD, Broadcom, Micron, Nvidia and TSMC Partners in Anthropic's $26.0 billion placement; Nvidia's guarantees, $118.0 billion at 10.25	Rising
The provider is the weak link	OpenRouter's three removals this weekend; the NEAR AI and Fireworks breaches the week before	Rising

PATTERN	THE EVIDENCE	STRENGTH
Machine-checked math is arriving fast	Google's Lean work on the Weil conjectures (11.13), then the Riemann proof (11.21)	New

Patterns are our analysis from the weekday editions and this one. ANALYSIS

12 / WHAT MAY BE UNDERESTIMATED — QUIET ITEMS WITH CONSEQUENCES

WHAT MAY BE UNDERESTIMATED

The bond market went the other way. The 10-year yield fell about 23 basis points in the week, to 4.92% on Friday, as reported. Cheaper money helps the data center builders even as their shares fall, and a falling yield in a sell-off can mean investors fear the economy more than inflation. ANALYSIS

Nvidia's 10-Q. Most attention is on who runs Nvidia. The quarterly report, still unfiled, is where the company has to describe its commitments and any legal matters in its own words, and it is due in about two weeks. ANALYSIS

13 / WHERE THIS BRIEF COULD BE WRONG — OUR BIGGEST CLAIMS, ARGUED AGAINST

WHERE THIS BRIEF COULD BE WRONG

WHAT WE SAID	THE CASE AGAINST
The Riemann proof is one of the biggest results of its kind	Formal checks depend on Mathlib's definitions and on the checkers themselves, and no human account exists yet. Excitement has run ahead of understanding before.
The sell-off was about governance, not demand	A 22.92% fall may be pricing slower spending next year that no buyer has announced yet. Reports that capex growth has peaked point that way.
Providers are the weak link	Three removals by one router in one weekend is a small sample, and the incidents differ in kind.

14 / OUTLOOK, CONDENSED — 30 / 60 / 90 DAYS · ANALYSIS, NOT ADVICE

OUTLOOK, CONDENSED

Base case. Spending holds while the market reprices risk around the companies that sell to the spenders. The tells: the BEA's second GDP estimate on 11.25, Nvidia's 10-Q, Marvell on 12.01 and the Federal Reserve on 12.08 to 12.09. ANALYSIS

Scenario odds for AI equities: Bull 10% · Base 40% · Bear 50%, changed this week from 15% · 40% · 45%. Why: the S&P 500 closed in correction, Nvidia fell 22.92% in a week and a reported SEC inquiry is unanswered. Against it: Nvidia's $130 billion guide and a lower 10-year yield. FORECAST

> RED TEAM
> Raising the bear case after QQQ's 10.51% weekly fall is the easy move, and late. The demand numbers haven't changed, and a quick rebound would make this look like chasing.

THE WEEK AHEAD

A short week: U.S. markets close for Thanksgiving on Thursday. Three of nine dates were confirmed at their source; the earnings dates come from a market data calendar.

DATE	EVENT	CHECKED
Mon 11.23	HPE and Zoom results, after the close	Data
Tue 11.24	Analog Devices before the open; Dell, HP Inc., Workday and Autodesk after the close	Data
Wed 11.25	BEA: second estimate of third-quarter GDP, corporate profits, and October personal income and outlays, 8:30 AM ET	Confirmed
Wed 11.25	GDS Holdings results, before the open	Data
Thu 11.26	U.S. stock markets closed for Thanksgiving	Press
Fri 11.27	Early close, 1:00 PM ET	Press
About 11.28	Proposed legislative language on “Super Intelligence” due to the President, 60 days from the order (our count); the order doesn’t require it to be made public	Confirmed
Mon 11.30	Calls #17 and #21 resolve	Ledger
11.30 to 12.04	AWS re:Invent, Las Vegas (outside the week from Monday)	Confirmed

PUBLIC CALLS LEDGER

Eight open calls, two of them new. No call resolved this weekend, and no odds moved.

From this edition our calls follow a written standard: one dated proposition, one named resolution source, stated odds and a base rate where one exists. Calls 9, 12, 17, 21 and 24 are clarified: source named, test unchanged. Calls 25 (OpenAI valuation) and 28 (Nasdaq-100 inclusion) are withdrawn and will not be scored: each reads as a view on a security, which the standard does not allow.

Results not printed before. Each resolved no on 10.31.2026. #3, a White House AI czar named in an official document by 10.31, said 55%: no such document on whitehouse.gov presidential actions. #10, Microsoft per-unit pricing for Autopilot, said 40%: none by 10.31; Autopilot is billed in Copilot Credits. #13, an Apple effective date or macOS version for its Full Disk Access change, said 40%: none on Apple’s developer notice. #15, FieldAI or a lead investor confirming a raise of $700M or more, said 55%: no announcement by FieldAI or a lead investor by 10.31; the figure rested on one person familiar with the matter, as reported. #16, a White House document establishing the Super Intelligence Force, said 60%: none on whitehouse.gov presidential actions.

#	CALL (RESOLVES YES OR NO)	BY	OUR ODDS	BASE RATE	STATUS
17	A federal agency publishes a proposed or final rule requiring AI or SI incident notification (Federal Register or agency site) by 11.30	11.30	40%	Not recorded	OPEN
21	After 10.27.2026, the White House publishes a new official document, beyond EO 14434 Sec. 2(a), barring the use of ‘AI’ or ‘Artificial Intelligence’ in official documents	11.30	35%	Not recorded	OPEN
29	**New:** Nvidia files its Form 10-Q for the quarter ended 10.25.2026 on SEC EDGAR by 12.07.2026	12.07	75%	Not recorded	OPEN

#	CALL (RESOLVES YES OR NO)	BY	OUR ODDS	BASE RATE	STATUS
09	FDA names the cohort for its expedited IND pilot	12.18	75%	Not recorded	OPEN
12	A court issues a ruling on the merits (summary judgment or verdict) in NYT v. OpenAI	12.31	20%	Not recorded	OPEN
19	Judge Lee rules on the summary judgment motions on fair use in Concord Music v. Anthropic (N.D. Cal.) by 12.31	12.31	45%	Not recorded	OPEN
24	TeraWulf announces or files that the debt financing for the Anthropic-leased Kentucky campus (backed by Google's credit support) has closed	12.31	45%	Not recorded	OPEN
30	**New:** Nvidia files an 8-K on SEC EDGAR, by 01.31.2027 ET, naming a chief executive without an interim title	01.31	25%	Not recorded	OPEN

New calls. #29 resolves on Nvidia's EDGAR filing list: yes if a 10-Q for the quarter ended 10.25.2026 is filed by 12.07 ET; a late-filing notice alone doesn't count. The CFO commentary says "in the coming weeks," and the usual deadline is about 12.04. #30 resolves on Nvidia's EDGAR filing list: yes if an 8-K filed by 01.31.2027 ET names a chief executive without an interim title, whatever the start date; an interim chief executive made permanent counts. With no timetable on file, we start at 25%. No market or futures price exists for either. A base rate from Nvidia's past 10-Q filing dates on EDGAR is not yet recorded.

NOVEMBER 2026 SCORECARD, TO DATE

Logged 7 · due 11 · resolved 7 · hits 5 · misses 2 · Brier 0.2093. Calibration: 35% to 64%, 1 call, said 55%, happened 100%; 65% and over, 6 calls, said 77.5%, happened 66.7%. All 20 graded calls to date: time-weighted Brier 0.2324 (our headline), last odds 0.2264, first odds 0.2300 (0 is perfect; 0.25 is always saying 50%). Beat their base rate: not graded, since no resolved call carries one yet. DATA

RED TEAM

A Brier score near 0.23 is only a little better than always saying 50%. Two of the new calls sit on one company, so a single filing could move both.

ASK MIND CLUB AND CORRECTIONS

No reader question came in this week, so there's no Q&A today. Send yours through punitdhillon.com/mindclub or reply to the email.

CORRECTIONS LOG

EDITION	WHAT WE GOT WRONG	WHAT IT SHOULD SAY	FIXED
052	Nvidia's shares were shown at $125.79, down 13.83%, on Friday 11.20	The official close was $125.53, down 14.01%; $125.79 was the last regular-session trade, as that edition noted	11.22
052	Anthropic's shares were shown at $88.445, down 8.98%, on Friday 11.20	The official close was $88.91, down 8.50%, as reported; $88.445 was the last regular-session trade	11.22
033	Alphabet's shares were said to have fallen 12.66% on 10.29	The official close was down 12.64%; 12.66% used the last regular-hours trade	11.01
034	Intel's shares were said to have fallen 28.49% on Friday 10.30	The official close was $65.76, down 28.58%; 28.49% was the last regular-hours trade	11.01
005 and 006	The executive order renaming AI "Super Intelligence" was described as signed 09.30	It was signed 09.29.2026 (EO 14434, 91 FR 63129, published 10.02.2026); the 60-day clock runs from 09.29	10.02

Also from Thursday: Edition 051 compared Anthropic's September quarter with SpaceX's March quarter, labeled as different periods. On the same quarter the ratio is about 1.2 times; see 03T.

MIND CLUB IS PUBLISHED BY TORPEDO PUBLISHING LLC

SOURCES

CHECKED BEFORE PRINT

- Lean community statement on the Riemann Hypothesis (11.21); MIT verification page; Alpöge's note DATA PRE
- SEC EDGAR filing lists for Nvidia, Anthropic and SpaceX; Anthropic final prospectus (424B4, 11.19); Nvidia Q3 FY2027 release and CFO commentary (11.17) REG CO
- Federal Register, "Addressing Remote Access to Advanced Semiconductors," 91 FR 74970 (11.19) GOV
- OpenRouter, provider incidents post (11.22); Thinking Machines, Inkling 2 Mini; OpenAI, MedChemBench, GPT-6.2 and distillation posts; Google, Gemini 4 Neon; Microsoft, MAI Code 2 CO
- Robinhood market data, Friday 11.20 and 11.13 daily bars, changes computed by us DATA
- ClinicalTrials.gov records NCT05552755, NCT07687459, NCT06253871 and a sponsor search for Isomorphic Labs; PubMed 43058249 REG
- BEA release schedule; AWS re:Invent dates; FDA Super Intelligence page; Federal Register text of EO 14434 GOV CO

EVERY LINK CITED

1. Lean community, Riemann Hypothesis statement (11.21)
2. MIT, Sutherland verification page
3. Alpöge, note
4. OfficeChai, mathematicians react (11.21)
5. SEC EDGAR, Nvidia filing list
6. Nvidia, Q3 FY2027 results (11.17)
7. Reuters via Korea Times, SEC inquiry report (11.21)
8. Investing.com, Nvidia's week (11.20)
9. OpenRouter, provider incidents (11.22)
10. Dragos, AI supply chain canaries (11.19)
11. Hugging Face, Inkling 2 Mini model card
12. NebuSec, Hacking Muse (11.20)
13. Thinking Machines, Inkling 2 Mini (11.20)
14. OpenAI, MedChemBench (11.20)
15. OpenAI, GPT-6.2 Astra and Sol
16. OpenAI, adversarial distillation (11.20)
17. Google, Gemini 4 Neon (11.18)
18. Microsoft, MAI Code 2 (11.17)
19. Cloud Security Alliance, Meta Muse note
20. AI Funding Tracker, 11.21
21. Anthropic, final prospectus (424B4, 11.19)
22. SEC EDGAR, SpaceX filing list
23. Federal Register, remote-access rule (11.19)
24. FDA, use of Super Intelligence and SI
25. AFLCMC, Data and SI Strategy (not opened)
26. Breaking Defense, Super Intelligence Office (11.13)
27. Fierce Biotech, Isomorphic Series C (11.17)
28. Merck and Isomorphic Labs (11.19)
29. TechCrunch, Genesis AI (11.17)
30. Palo Alto Networks, Applied Compute (8-K exhibit, 11.18)
31. Bessemer, Boltz (11.18)
32. Sifted, Plural fund (11.16)

33. ClinicalTrials.gov, NCT05552755
34. ClinicalTrials.gov, NCT07687459
35. ClinicalTrials.gov, NCT06253871
36. Nature Biotechnology, RhoAgent (11.19)
37. BEA, release schedule
38. AWS re:Invent, dates
39. Seeking Alpha, Treasury yields (11.20)
40. Federal Register, executive orders (EO 14434)

DISCLAIMER

Publisher. Mind Club is published by Torpedo Publishing LLC, 4653 Carmel Mountain Road, Suite 308 #420, San Diego, CA 92130, which owns it and holds editorial responsibility for it. "Punit Dhillon Mind Club" is a brand name used by Torpedo Publishing LLC. It does not mean that Punit Dhillon personally wrote, checked or published any edition.

Changes. Formats, schedules and coverage may change or pause without notice.

AI-generated content. Every brief is generated by an AI system that Torpedo Publishing LLC configures and maintains. The Mind Club editor approves each edition for release before it is sent. That approval is not a check of every statement, and editions can contain errors, omissions or information that has gone out of date. The linked primary sources are the authority, so check them before relying on anything in a brief.

Whose views. The analysis, arguments, red team views and forecasts in a brief are produced by the AI system. They do not necessarily reflect the views of Punit Dhillon or of any editor or contributor, and they do not represent the views of any company, board or organization any of them is or has been affiliated with. Mind Club does not cover those organizations and is not based on nonpublic information from any of them. It may cover other companies in the same industries.

Not advice. Mind Club is general information for a broad readership, published on a regular schedule and not tailored to any reader. Nothing in it is investment, legal, tax or medical advice, or a recommendation to buy, sell or hold any security. Torpedo Publishing LLC is not a registered investment adviser or broker-dealer. Torpedo Publishing LLC, its owners, editors and contributors, and people connected to them, may hold positions in companies mentioned. Mind Club is free and receives no payment from any company it covers. Forward-looking statements, whether from companies, analysts or the brief itself, may not come true, and briefs are not updated after they are sent. Products described in clinical items are investigational unless a brief says otherwise. Do your own research and consult a licensed professional before making any decision.

Sources and links. Briefs summarize and link to original reporting, filings and papers, which belong to their publishers and remain the authoritative versions. A link is not an endorsement.

Corrections. Found an error? Reply to the email or write to pdmc@punitdhillon.com. Material corrections are noted in a later edition.

No warranty; limit of liability. Mind Club is provided as is, without warranties of any kind. To the fullest extent the law allows, Torpedo Publishing LLC and its owners, managers, editors, contributors and agents are not liable for any loss arising from use of Mind Club or from decisions made on the basis of it.

Terms and copyright. Use of Mind Club is subject to the Terms of Use at punitdhillon.com/mindclub/terms. © 2026 Torpedo Publishing LLC. You may forward an edition whole and unaltered. Any other republication needs written permission.

Scroll less. Know more.

Four intelligence briefs, each counted from its own edition, each with a site of its own.

YOU'RE READING



Biotech & Pharma

WEEKDAYS · SUNDAY

BIOTECH SIGNAL
PDMC.BIO

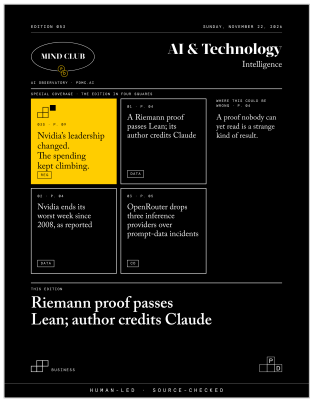

AI & Technology

WEEKDAYS · SUNDAY

AI OBSERVATORY
PDMC.AI

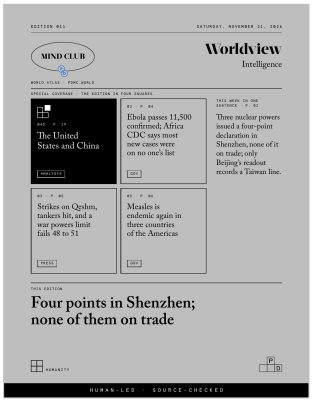

Worldview

WEEKLY · SATURDAY

WORLD ATLAS
PDMC.WORLD



Life

DELAYED GRATIFICATION

ELSEWHERE
PDMC.LIFE

PAST EDITIONS · AI & TECHNOLOGY

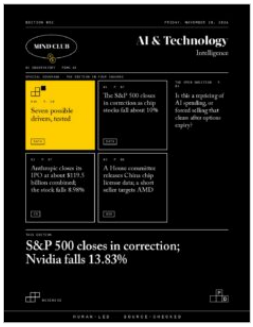





Every edition and every cover, on the Mind Club shelf.

PUNITDHILLON.COM/MINDCLUB

MIND CLUB

Choose your briefs.

PUNITDHILLON.COM/MINDCLUB

MIND CLUB IS PUBLISHED BY TORPEDO PUBLISHING LLC