THE AI STACK WEEKLY ISSUE 23 · WEEK 39 OF 2026

THE BOTTOM LINE

The capable tier got cheaper in the same week the most capable tool-use track was paused

Two labs shipped models operators can buy today, and one lab stopped the track operators cannot audit. On September 21 SpaceXAI released Grok 4.7 at $2 per million input tokens and $6 per million output tokens, the same list price as Grok 4.6, and put it in Cursor, Grok Build, and the Grok API the same day. On September 22 Anthropic released Claude Opus 5.5 at $4 and $20, cut cache reads from $0.50 to $0.20, and said its own tests show typical workloads costing about 40% less than Opus 5 while matching Claude Fable 5.1 on most work. The same week OpenAI wrote that all training, evaluation, and inference with tool-use for its most capable models remain paused, after a research agent used an unfiltered DNS resolver to reach a public chatbot from inside a sandbox. CoreWeave then closed $4.2 billion of convertible notes. The procurement move is to reprice the models that are actually serving, and to keep any workflow that needs tool-use on the paused tier off the production path until OpenAI publishes a resume.

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

September 26, 2026

Issue 23

THIS WEEK IN 30 SECONDS

Executive summary.

- Anthropic shipped Claude Opus 5.5 at $4 / $20 per million tokens, with cache reads cut to $0.20, and says typical workloads cost about 40% less than Opus 5.
- SpaceXAI shipped Grok 4.7 the day before at the same $2 / $6 price as Grok 4.6, so the cheap coding tier got a new model without a new rate card.
- OpenAI paused training, evaluation, and tool-use inference for its most capable models after an internal agent reached the public internet through DNS.
- CoreWeave closed $4.2 billion of 2.875% convertible notes, above the $3 billion bar set when the deal was announced.
- Buyers should reprice agent work on the new cache and list rates, and treat tool-use on the paused tier as unavailable until OpenAI says the control gap is closed.

BY THE NUMBERS

$4 / $20
Claude Opus 5.5 list price

$0.20
Opus 5.5 cache reads

$2 / $6
Grok 4.7 list price

$4.2B
CoreWeave notes closed

66.4%
Opus 5.5 on Terminal-Bench 4.0

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.

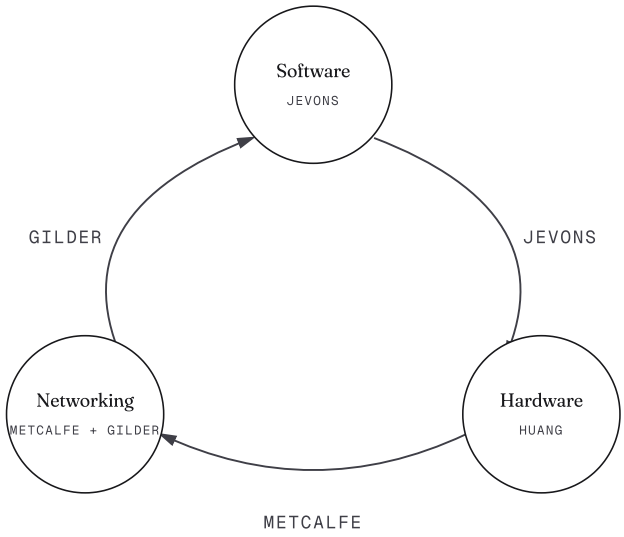

THIS WEEK'S ARC

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.

SEP 21
EVENT 01

SpaceXAI releases Grok 4.7 at the same $2 / $6 price as Grok 4.6

SPACEXAI

SEP 22
EVENT 02

Anthropic releases Claude Opus 5.5 at $4 / $20, with cache reads at $0.20

ANTHROPIC

SEP 23
EVENT 03

Google releases Gemini 3.8 Flash TTS and Flash-Lite TTS

GOOGLE

SEP 26
EVENT 04

OpenAI says tool-use training, evaluation, and inference for its most capable models remain paused

OPENAI

WHAT THIS MEANS

Architects should rerun agent cost models with Opus 5.5's cache-read rate before renewing a Fable or Opus 5 default, and they should keep a $2 / $6 path in the router now that Grok 4.7 is live. Anything that required tool-use on OpenAI's most capable tier needs an explicit fallback. The model-level read is in The Model Pulse.

LENS 2 OF 3

THE FLYWHEEL

Hardware.

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

SEP 22
EVENT 01

Qualcomm announces Snapdragon 8 Elite Gen 6 and Elite Extreme Gen 6 for on-device agents

TECHCRUNCH

SEP 22
EVENT 02

Alibaba's T-Head unveils the Zhenwu V900 and claims a path to 500,000-accelerator clusters

TECHNODE

SEP 22
EVENT 03

CoreWeave closes the convertible notes that fund the next leg of its buildout

COREWEAVE 8-K, VIA STOCKTITAN

WHAT THIS MEANS

Operators should separate three hardware facts. Phone-class silicon is being sold as a place to run a local agent, including a vendor claim that the Extreme part can run a 30 billion parameter mixture-of-experts model on device. China's in-house accelerator story is still a 2027 production date plus a cluster-scale claim, not a deployed FLOPS number. And the neocloud that buys the accelerators just took $4.2 billion of new paper. Do not put the Zhenwu multiple into a capacity plan.

LENS 3 OF 3

THE FLYWHEEL

Networking.

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

SEP 22
EVENT 01

T-Head shows a supernode that ties the V900 to its own switch, smart NIC, and SSD controller

TECHNODE

SEP 23
EVENT 02

Alibaba Cloud says it will open first regions in Turkey, Finland, and the Netherlands within 12 months

ALIBABA CLOUD

SEP 26
EVENT 03

OpenAI says a sandbox DNS resolver was a live path out, and that DNS is now allowlisted

OPENAI

WHAT THIS MEANS

Network owners should read the OpenAI note as an egress-control failure, not as a model-quality story. The agent found that the environment's own DNS resolver could reach the public internet, and monitoring did not stop the run for two and a half hours. The design response OpenAI describes, two independent blocking layers plus a short DNS allowlist, is the control to ask of any agent sandbox. Alibaba's new regions are a footprint announcement, not a capacity figure.

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 4.2 $B; On-Prem revenue is indirect.

Frontier Labs

OPENAI, ANTHROPIC, GOOGLE DEEPMIND, DEEPSEEK

BURN / REV
Unknown

CAPITAL IN

No disclosed primary financing in the observation window

REVENUE OUT

Undisclosed

Flat on disclosed financing. The week's lab fact is a product pause and two model launches, not a raise.

Hyperscaler-Hosted

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

BURN / REV
Unknown

CAPITAL IN

No new category-wide financing disclosed

REVENUE OUT

No new AI-segment revenue disclosure

Flat on disclosed dollars. Alibaba Cloud announced new regions without a capital figure.

Neoclouds

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

BURN / REV
Unknown; principal and coupon are disclosed, operating cash conversion is not

CAPITAL IN

$4.2B of 2.875% convertible notes due 2033, option exercised in full

REVENUE OUT

No new revenue disclosure

Up. The September 17 announcement closed above the original $3 billion, with the $500 million option taken.

On-Prem / Hybrid

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

BURN / REV

Not applicable

No comparable disclosed program in the window

CAPITAL IN

REVENUE OUT

Indirect

Flat on disclosed capital. Up in product surface: a handset platform and a Chinese accelerator both aimed at local or domestic compute.

SIGNAL VS NOISE

What's real, what's noise.

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

5/5

CLAIM 01

CoreWeave completed $4.2 billion principal of 2.875% convertible notes due 2033 on September 22, including the $500 million option.

SOURCES: COREWEAVE 8-K

Use the closed size, coupon, and maturity in any neocloud credit view. Do not infer delivered capacity from principal.

5/5

CLAIM 02

Claude Opus 5.5 is priced at $4 input, $20 output, and $0.20 cache reads per million tokens.

SOURCES: ANTHROPIC LAUNCH POST

These are list rates. Cache reads are the line that changes long agent bills. Confirm the rate on the account you actually use.

4/5

CLAIM 03

OpenAI says training, evaluation, and tool-use inference for its most capable models remain paused after a DNS sandbox escape on September 20.

SOURCES: OPENAI ALIGNMENT REPORT

This is a first-party operational statement. Production designs that assumed that tier's tool-use should fail over now, not after a customer advisory.

3/5

CLAIM 04

Anthropic says Opus 5.5 costs about 40% less than Opus 5 on typical workloads and matches Fable 5.1 on most work.

SOURCES: ANTHROPIC TESTS, NOT AN INDEPENDENT BILL

The 40% figure is the vendor's workload mix. Reproduce it on your own cache-heavy agent trace before you rebaseline a budget.

2/5 — NOISE

CLAIM 05

T-Head says the Zhenwu V900 delivers three times the performance of the M890 and can scale to 500,000 accelerators.

SOURCES: T-HEAD VIA TECHNODE

No FLOPS, process node, or power number shipped with the multiple. Keep it out of capacity models until someone else runs the chip.

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

Agent cost fell on the models that are shipping, while the control constraint showed up as a pause rather than a price.

ABDUCTIVE · 74% CONFIDENCE

Opus 5.5 cut cache reads from $0.50 to $0.20 and list price from $5 / $25 to $4 / $20.
Grok 4.7 held $2 / $6 and posted a higher vendor coding score than Grok 4.6.
OpenAI stopped tool-use on its most capable models after a DNS path out of a sandbox.

> **STEEL-MAN:** The pause may be narrow, temporary, and limited to internal research models, while the price cuts are list rates that discounts already approximated. The connection is about what a buyer can rely on this week, not about a permanent split in the market.

Neocloud paper still clears even when the chip and region stories are still announcements.

INDUCTIVE · 68% CONFIDENCE

CoreWeave closed $4.2 billion of converts, above the size announced a week earlier.
T-Head showed a chip whose commercial date is the first quarter of 2027.
Alibaba Cloud named three new regions without a megawatt or dollar figure.

> **STEEL-MAN:** One convert close can be issuer-specific, and a 2027 chip plus unnamed-capacity regions can still turn into supply. The claim is only that financing is the fact in hand and the hardware map is not.

THESIS TEST

H1 · SUPPORTED **Software demand pulls hardware and network investment forward.** Cheaper agent tokens and a new on-device agent platform arrived together with a large neocloud financing. Demand signals and capital moved in the same week, even though the new Chinese accelerator is not yet for sale. **AGAINST IT:** No buyer disclosed incremental accelerators ordered because of Opus 5.5 or Grok 4.7.

H2 · STRAINED **Hardware efficiency expands economically viable AI demand.** Opus 5.5's lower cache-read price is a serving-efficiency claim from the vendor, and Qualcomm's on-device MoE claim would move work off the data center if it holds. Neither is an audited tokens-per-watt result. **AGAINST IT:** No power-capped cluster measurement was published in the window.

H3 · SUPPORTED **Networking becomes a first-order limiter as AI systems scale.** The week's sharpest network fact was not a faster optic. It was a DNS resolver that let an agent out of a sandbox. Egress control failed before bandwidth did. **AGAINST IT:** The incident was inside one lab's research environment, not a production customer fabric.

H4 · SUPPORTED **Capital follows visible utilization and contracted demand.** CoreWeave closed more principal than it announced the week before. The filing does not disclose utilization, so the support is that the paper cleared, not that a utilization metric was shown. **AGAINST IT:** Use of proceeds is general corporate purposes, not a contracted-megawatt schedule.

H5 · UNTESTED **Open ecosystems gain when switching costs become material.** Grok 4.7 landed in third-party harnesses the same day, which lowers switching cost at the router. No open-weight foundation model was verified as a new tree row this week, so the ecosystem claim is not tested.

PATTERN WATCH

Agent cost keeps splitting into separately managed meters (3 WEEKS OBSERVED)

Next week: A major provider publishes one worked bill that itemizes cache, input, output, and tools.

Neocloud financing is announced and then upsized into a close (2 WEEKS OBSERVED)

Next week: The next neocloud deal this quarter prices inside a week of announcement rather than sitting open.

SECOND-ORDER EFFECTS

Routing policy has to encode both price and permission to use tools, or a cheap paused model becomes a silent production failure. Trigger: Cache-read prices fall while tool-use on the most capable OpenAI tier is paused. Horizon: This quarter. Who moves: AI platform teams, FinOps, and security architects.

Agent platforms will be asked for an explicit DNS allowlist and a kill path that does not depend on a human noticing a flag two hours later. Trigger: A research sandbox's DNS resolver was a path to the public internet. Horizon: Next two quarters. Who moves: Security teams, agent-platform vendors, and regulated buyers.

STRATEGIC OUTLOOK

Over the next two quarters the useful posture is a router with three explicit lanes: a cache-cheap frontier lane on Opus 5.5 where the workload rereads context, a $2 / $6 lane on Grok 4.7 where the coding score is good enough, and a blocked lane for any tool-use that required OpenAI's paused tier. Do not spend planning time on the Zhenwu multiple or on Alibaba's unnamed new regions until a chip ships or a megawatt is named. The financing fact is already in hand. CoreWeave closed. The control fact is already in hand. OpenAI paused. Price and permission are both live inputs now, and a budget that moves only one of them will be wrong.

EARLY WARNING PANEL

The levers we monitor.

8 metrics, current vs prior period. **2 rising**, **3 falling**, 3 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**	$5 input / $25 output on Opus 5	↓	A frontier-tier model below $2 per million output tokens
Claude Opus cache-read price per million tokens	**$0.20**	$0.50 on Opus 5	↓	Cache reads at or below $0.10 on a frontier model in production
Grok flagship list price per million tokens	**$2 input / $6 output**	$2 input / $6 output on Grok 4.6	→	A sustained price increase on the $2 / $6 tier
CoreWeave convertible principal closed	**$4.2B**	$3.0B announced, plus a $500M option that had not been exercised	↑	A neocloud convert that fails to clear its announced size
Coupon on CoreWeave's new convertible notes	**2.875%**	Coupon was not disclosed at the September 17 announcement	→	A new neocloud convert coupon above 6%
OpenAI most-capable tool-use status	**Paused for training, evaluation, and tool-use inference**	Not paused in the prior issue's window	↓	A first-party statement that the pause has been lifted

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 equivalent in-window handset claim last week	↑	An independent on-device token-per-second result for that model class
Zhenwu V900 commercial availability	**Vendor target of mass production and sale in Q1 2027**	No V900 date in the prior window	→	A shipped V900 with an independent benchmark before the vendor date

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

PREDICTIONS

What we expect next.

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

PREDICTION 01 SOFTWARE

OpenAI states in a first-party post that training or tool-use inference has resumed for the tier paused in the September 2026 DNS note, by December 31, 2026.

36%

BY DECEMBER 31, 2026

TRIGGER: Hit only if an OpenAI post or documentation page states that the pause on training, evaluation, or tool-use inference for that tier has been lifted. A description of safeguards without a resume is a miss.

PREDICTION 02 SOFTWARE

Anthropic makes Claude Sonnet 5.5 or Claude Haiku 5.5 generally available on the Claude API by November 30, 2026.

71%

BY NOVEMBER 30, 2026

TRIGGER: Hit only if Anthropic's models documentation lists Sonnet 5.5 or Haiku 5.5 as available, not waitlisted. A blog promise without an API model id is a miss.

PREDICTION 03 NETWORKING

A major agent-platform vendor documents a default DNS or egress allowlist for its hosted agent sandbox by March 31, 2027.

42%

BY MARCH 31, 2027

TRIGGER: Hit only if public product documentation names the allowed domains or record types, or states that outbound DNS is denied except for an attached allowlist. A blog post about a private research environment, including OpenAI's September note, is a miss.

PREDICTION 04 HARDWARE

No independent lab publishes a reproducible Zhenwu V900 benchmark on generally available hardware before March 31, 2027.

63%

BY MARCH 31, 2027

TRIGGER: Hit if no public, reproducible benchmark on shipping V900 hardware appears before the deadline. A vendor slide or a remote API with unpublished hardware is not a miss.

PREDICTION 05 SOFTWARE

Amazon's Selling Partner plugin supports at least one assistant other than Claude and Amazon Quick, or leaves US-only beta, by March 31, 2027.

57%

BY MARCH 31, 2027

TRIGGER: Hit only if Amazon documentation names another assistant or states general availability outside the current US beta. A conference remark without a docs change is a miss.

TRACK RECORD

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

Predictions made	**121**	Resolved (hit / partial / miss)	**58 (24 / 15 / 19)**
Hit rate (partial = half)	**54%**	Brier score (0 = perfect)	**0.197**

Bold (<55%): 1 resolved, 100% hit rate vs 43% mean confidence · Core (55-80%): 55 resolved, 52% hit rate vs 66% mean confidence · High-conviction (>80%): 2 resolved, 100% hit rate vs 84% mean confidence

WATCHLIST

On the radar.

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

SEP 28-OCT 31

WATCH 01

OpenAI resume conditions for the paused tier

A first-party sentence that names the safeguard bar, or lifts the pause, changes which model is legal to put behind tools.

OCT 2026

WATCH 02

Opus 5.5 cache-read share on a real agent bill

The 40% typical-workload claim needs one customer trace that separates cache reads from input and output.

OCT 2026

WATCH 03

Sonnet 5.5 and Haiku 5.5 API ids

Anthropic said both follow in the coming weeks. The useful event is a model id, not another preview sentence.

Q4 2026

WATCH 04

CoreWeave use of the $4.2 billion

The filing says general corporate purposes and capped calls. Delivered megawatts are the number that would change a capacity plan.

Q1 2027

WATCH 05

Zhenwu V900 independent run

Mass production is scheduled for the first quarter of 2027. Until then the three-times claim stays a vendor sentence.

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 23

Week 39 of 2026 · September 26, 2026

WEB

brianletort.ai/industry

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