

AI Weekly: Gate week — Astra pause, Mythos scans, Sol \$4/\$20

Weekly Digest · August 22, 2026 · 8 min · August 17, 2026 – August 22, 2026

Between 17 and 22 August 2026 the labs gated access, paused scaling, and repriced. Not a chatbot circus. OpenAI put a two-week RL pause on deployment-bound training and flagged Astra's cyber bar. Anthropic showed wet-lab protein binders, then GA'd computer use with Skills and Files. Mythos 5 landed as a scanner inside Claude Security. Sol got a promotional \$4/\$20. DeepSeek flipped V4 to peak and off-peak rates. Each section: the idea, then the claim. Sources at the bottom.

1. OpenAI: two-week RL pause, Astra at Critical cyber

Frontier models get better by training harder. This week OpenAI said some of that training pauses — on purpose — while cyber risk is watched.

OpenAI announced a two-week reinforcement-learning pause on the latest models intended for deployment. The largest planned frontier RL run stays on hold. Smaller-scale training and evals continue. A Hugging Face eval incident, plus Astra, may meet a Critical cybersecurity threshold (determined 7 August). Monitoring is required for RL-plus-tools at Sol capability or higher. Astra-plus-tools inference gets extra monitoring — about 20% of watched inference compute. Some Astra workloads remain paused until they migrate.


So why does it matter? The gate is on training and monitored inference, not on turning ChatGPT off.

Two-week RL pause on deployment-bound models. Not a Sol kill switch.

Largest frontier RL still on hold; smaller work continues.

Astra may hit Critical cyber; ~20% of watched inference is under that monitor.

OpenAI — pacing / Astra



Pause · Hold · Astra Critical possible (7 Aug) · ~20% watched inference. See Sources below. [Source](#)

2. Claude protein binders, wet lab

Protein design is usually a slog of failed campaigns. Anthropic's claim this week is wet-lab hit rates that beat the usual 10–15% band — with Claude in the loop, not instead of the lab.

Mythos Preview and Opus 4.8 designed against 15 targets. Binders landed against 14. In a multi-target 48-hour window: Opus 4.8 at 22.6%, Mythos Preview at 26.7%. Mythos single-target hit 35.1%. Typical campaign range is 10–15%. Wet lab ran with Adaptiv Bio and Twist. Opus 5 timed NMR at 23 minutes and LC-MS at 19 minutes; purity 96.4% versus lab 96.33%. Life-science work is still blocked on the most capable model; an access program is promised, not shipped here.

So why does it matter? The story is a lab campaign with named hit rates — not “Claude replaced protein engineers.”


14 of 15 targets got binders; typical campaigns sit at 10–15%.

Multi: Opus 22.6%, Mythos 26.7%; Mythos single 35.1%.

Most capable model still blocked for life science in this post.

Claude — protein design hit rate

typical campaign 10–15% · Anthropic + Adaptyv/Twist



Opus multi 22.6% · Mythos multi 26.7% · Mythos single 35.1% · typical 10–15%. See Sources below. [Source](#)

3. Computer use + Skills + Files GA

Agents that click around a desktop used to be preview theater. On 20 August Anthropic called computer use, Skills, and Files generally available on the Claude Platform.

GA includes multi-action turns and browser use that sees page structure, not pixels only. HIPAA BAA coverage is in the post. Skills API lets orgs upload and version skills. Files API adds auto expiry, 5x rate limits, and 1 TB per org. Skills and Files land on Microsoft Foundry; computer and browser use are coming to Vertex. A named customer quote says 32 minutes down to 13 and about 30% cheaper — that is a vendor-cited customer story, not an audit.


So why does it matter? The product surface is GA plumbing — skills, files, and a desktop/browser agent — not a demo reel.

Computer use GA 20 Aug; multi-action turns; HIPAA BAA.

Files: auto expiry, 5x rate limits, 1 TB/org; Skills/Files on Foundry.

32→13 min / ~30% cheaper is a customer quote. Not an independent redo.

Claude Platform — computer use GA



Computer use · browser use · Skills API · Files API (5x limits · 1 TB/org). See Sources below. [Source](#)

4. Mythos 5 in Claude Security

Cyber-capable models keep getting talked about as if they were ChatGPT with a red badge. Anthropic's Mythos 5 story this week is a scan product, not a prompt box.

Mythos 5 ships inside Claude Security as an Enterprise public beta. Billing is standard tokens — no add-on. Findings come back as CWE, severity, confidence, and a suggested fix. A human must approve patches. Interactive patching uses the org's existing models, not Mythos-as-chat. A \$35M Defender Advantage Fund (0xDAF) puts credits on the table. The Cyber Verification Program is expanding. This digest is the existence of that gate — not an exploit walkthrough.


So why does it matter? Defenders get a scanner with structured findings. They do not get Mythos in a public chat UI.

Scan, not chat. Enterprise public beta.

CWE + severity + confidence + suggested patch; human approves.

\$35M 0xDAF credits; no exploit how-to here.

Mythos 5 — scan not chat



Scan · CWE findings · Enterprise beta gate · \$35M 0xDAF. See Sources below. [Source](#)

5. GPT-5.6 Sol \$4/\$20 promo

Token price is the bill people fight about. On 21 August OpenAI cut Sol's promotional API rate — with an end date in the posts we cite.

GPT-5.6 Sol API promo is \$4 input / \$20 output per 1M tokens. That is –20% input and –33% output versus \$5/\$30. The promo runs at least through 21 November 2026. Plus and Pro subscription prices are unchanged in the posts we cite.

So why does it matter? Cheaper Sol tokens for a dated window. Not a forever list price rewrite.


\$4 in / \$20 out per 1M; –20% / –33% vs \$5/\$30.

Promo at least through 21 Nov 2026.

Plus/Pro subs unchanged here.

OpenAI Sol — \$/1M tokens (promo cut)

promo at least through 21 Nov 2026



Was \$5/\$30 · promo \$4/\$20 · at least through 21 Nov 2026. See Sources below. [Source](#)

6. DeepSeek V4 peak/off-peak live

Flat token rates are easy to quote. Peak and off-peak make the bill depend on when you call.

DeepSeek's V4 pricing went live 16 August 2026 at 16:00 UTC (00:00 Beijing on 17 August). Peak windows are 01:00–04:00 and 06:00–10:00 UTC. Off-peak is half of peak. Flash output moved from a flat \$0.28 to \$0.66 off-peak / \$1.32 peak. Pro output moved from \$0.87 to \$1.98 / \$3.96.

So why does it matter? V4 is not a small Flash nudge. The schedule is half the story.


Live 16 Aug 16:00 UTC; peak 01:00–04:00 and 06:00–10:00 UTC.

Off-peak = half of peak.

Flash out \$0.28 → \$0.66 / \$1.32; Pro out \$0.87 → \$1.98 / \$3.96.

DeepSeek V4 Flash — \$/1M output


V4 Flash output · live 16 Aug 16:00 UTC



V4 Flash out: old \$0.28 · off-peak \$0.66 · peak \$1.32 · live 16 Aug 16:00 UTC. See Sources below. [Source](#)

DeepSeek V4 Pro — \$/1M output

V4 Pro output



V4 Pro out: old \$0.87 · off-peak \$1.98 · peak \$3.96. See Sources below. [Source](#)

Frequently Asked Questions

Did OpenAI stop shipping GPT-5.6 Sol?

No. The pause is reinforcement learning on the latest deployment-bound models. ChatGPT/API Sol already out is not “turned off” in this post. Largest planned frontier RL remains on hold. Smaller-scale training continues.

Did Claude replace protein engineers?

No. Wet-lab hit rates vs a typical 10–15% campaign. 14 of 15 targets got binders. Life-science work is still blocked on the most capable model; an access program is promised, not shipped here.

Can anyone prompt Mythos 5 in a chat box?

No. Claude Security returns findings (CWE, severity, confidence, suggested patch). Interactive patching uses the org’s existing models. Enterprise public beta. Not Mythos-as-ChatGPT.

Is \$4/\$20 the permanent Sol list price?

No. Promotional at least through 21 November 2026. Input –20%, output –33% vs \$5/\$30. Plus/Pro subscription prices unchanged in the posts we cite.

Did DeepSeek only nudge Flash a little?

No. V4 moved to peak/off-peak. Off-peak is half of peak. Live 16 August 16:00 UTC. Flash output \$0.28 → \$0.66 / \$1.32. Pro output \$0.87 → \$1.98 / \$3.96.

Sources

This is original synthesis. Underlying facts and charts are drawn from the reporting and vendor posts below.

[🔗 OpenAI — Pacing model development](#)

[🔗 Anthropic — Claude protein design](#)

[🔗 Claude — Computer use, Skills API, Files API](#)

[🔗 Claude — Mythos 5 to more defenders](#)

[🔗 OpenAI API changelog — Sol pricing](#)

[↗ OpenAI — GPT-5.6 \(price update note\)](#)

[↗ OpenAI — API pricing](#)

[↗ DeepSeek — API changelog](#)

Vendor numbers, customer quotes, and first-party benches are not independent audits. Sol \$4/\$20 is a promo window through at least 21 November 2026. DeepSeek peak/off-peak hours are as published. Mythos 5 ships as a scanner inside Claude Security — not a public chat box. OpenAI's Astra-linked RL pause is not a product shutdown. Protein results are a lab campaign, not a drug. Educational synthesis only; no exploit how-to. Read the primary sources.