Open weights are a weapon now — one country funds them, another wants to gate them

July 14, 2026 · by Mark Bartlett

Download this post as PDF

InsiderLLM Weekly issue 13 – July 14, 2026

The models on your drive didn't change this week — what they mean did. Two stories landed that look unrelated and aren't: China is funding open weights like national infrastructure, and the West spent the same week arguing about whether to gate them. Plus the twist nobody saw coming — the country that made open weights a strategy might be first to lock down its own frontier models.

China Is Funding Open Weights Like Infrastructure

DeepSeek closed its first outside round in June — roughly $7.4 billion, at a $50B-plus valuation. One detail is worth reading twice: of every backer in that round, the only one that walked away with actual voting rights and no lock-up was China's state AI fund. Tencent, CATL, and Liang Wenfeng's own $3 billion all went into a locked partnership with no say. Then, weeks later, DeepSeek started shopping a second round at a reported $71B pre-money.

This is a company that gives its weights away for free, bankrolled like a national asset, with the state holding the wheel. That's not charity. It's policy.

The honest hedge: the $71B round is reported by the FT and The Information, not closed — treat it as direction, not done deal. But the direction is the point. Part of the reason a frontier-ish model fits on your 3090 is that a government decided open weights were worth this kind of money.

And here's the twist. The same country funding the giveaway is floating rules to restrict its own top-tier open models — a tiered proposal traced to a Supreme People's Court journal roundtable that would leave small tools on a simple filing but bar the most sensitive frontier models from public release. It's discussions, not law. But the tap that's open today has a hand near it.

📖 **The full breakdown — if China pulls back, who fills the gap — is here.**

Meanwhile, the West Is Arguing About Whether to Gate Them

Same week, Demis Hassabis dropped a manifesto calling for a US-led standards body — FINRA-style — that would screen frontier models before release, open or closed, whatever the country of origin. His warning is blunt: within 18 months, bio- and nuclear-relevant capability could sit inside open-source models that no government can pull back once they're out.

Then the Future of Life Institute's safety index landed and handed failing grades to three labs, one per continent — including Mistral (dead last) and DeepSeek. Mistral's response is the whole fight in a sentence: open weights are the safety mechanism. Independent scrutiny beats a handful of companies deciding behind closed doors what's safe for the rest of us.

Both can't be right, and the split lands on you specifically. One future screens the models before you ever get to download them. The other keeps the floodgates open and bets on daylight. I wrote up the honest case for the open side — including where it's genuinely underfunded, because pretending otherwise would be a vibe, not an argument.

📖 **The West's open-source second wind is here.**

Quick Hits

- **llama.cpp hit b10000.** Ten thousand builds. The engine most of us actually run crossed the milestone without so much as a press release. Pooling along, as usual.
- **DeepSeek V4 ships mid-July** with peak/valley API pricing and DSpark, their new speculative-decoding trick — 60–85% faster, no architecture change. Watch for the GGUFs to land, then we'll bench it.
- **GPT-5.6 dropped** with three tiers — Sol, Terra, Luna. Closed and cloud, so not our beat, but the Sol-vs-Fable coding-bench scrap is worth half an eye.

That's the week. The models on your drive didn't change — but the case for why they're there got a lot louder. Next edition lands next week.

— Mark, InsiderLLM

Running frontier-ish open weights on hardware they've got no business running on? Reply, or hit me at hello@insiderllm.com. I read everything.

Source: https://insiderllm.com/blog/newsletter-2026-07-14/

Free guides for running AI locally