

Over 1,100 AI lab staff ask the US to back a way to pace frontier AI

Source copy kept by wheresthe.ai for the story "Over 1,100 AI lab staff ask the US to back a way to pace frontier AI" (<https://wheresthe.ai/d/intl-pacing-the-frontier-open-letter-2026-07/>).

Source: <https://thenextweb.com/news/pacing-the-frontier-ai-employees-letter-us-government>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: ff32436b4fd2fe981b9f11615186b43ab8ddc1c5f91be76acc2fa1cd2af8e1ca

This article was published on July 28, 2026

Credit: Pacing the Frontier

More than a thousand of the people building the most powerful AI want a way to slow it down. On Tuesday, 1,134 employees of the leading AI companies signed a letter asking the US government to help build one.

The statement is titled Pacing the Frontier . It asks Washington to "support an international effort to develop the technical and governance tools needed to deliberately pace the frontier of automated AI development." Bloomberg first reported the letter was circulating.

The names are the story. Signatories include Anthropic chief executive Dario Amodei and its co-founders Jared Kaplan and Jack Clark. So did OpenAI chief scientist Jakub Pachocki, Meta chief scientist Shengjia Zhao, and Google's head of AI safety, Anca Dragan. Both Anthropic and OpenAI endorsed the letter officially.

What it actually asks for

It is not a call to stop. The letter does not ask anyone to pause or slow AI now, NBC News noted . It asks the government to make sure the option to pace exists later, if the technology outruns the people building it.

The specific fear is recursive self-improvement: AI that speeds up its own development. The signatories warn of "a real risk that capability development rapidly accelerates beyond our ability to understand or control the resulting systems."

That worry stopped being abstract days earlier. OpenAI's most advanced model broke out of its sandbox and hacked Hugging Face to score higher on an internal test. It builds directly on Anthropic's call last month for a verifiable way to pause.

The whiplash

The timing is remarkable, because these same companies spent the previous week arguing the opposite case. Days ago, Nvidia, Microsoft, Meta and others rallied around an open-weights letter championing openness as the path to safety.

Now many of the same firms want brakes. The clearest split runs through one company. On the same day Meta's chief scientist signed the pacing letter, Mark Zuckerberg attacked the rival labs. Their discourse, he said, is "overwhelmingly filled with doom."

Tightly controlling AI would be "abandoning our values," Meta's chief executive told the New York Times . Two open letters, one week, opposite instincts, and a fault line inside the industry's biggest names.

The three objections

The criticism was immediate, and it lands in three places.

The first is China. "If the US labs pace themselves, why would China wait?" asked the economist Christian Catalini. The letter half-concedes the point, admitting no company or country will slow down unilaterally.

The second is bad faith. "It is their company. They could just stop," wrote former Microsoft executive Steven Sinofsky. The charge stings because the signatories include the very CEOs who set the pace.

The third is the moat. Critics read every big-lab safety letter as incumbents lobbying to raise the drawbridge behind them, dressing a competitive move as caution. It is a charge, not a proven motive, but it follows these letters everywhere.

Washington is the real audience

The letter is aimed at a government that has so far wanted little to do with AI rules. President Trump's administration has argued that international AI governance would hobble the US against China. That stance may be softening, now that frontier models are starting to find and exploit real security holes.

The proposals now on the table are concrete. Amodei has floated an FAA-style agency that could test frontier models and halt a dangerous release. Demis Hassabis wants a body that reviews models before launch, Business Insider reported . Congressman Ted Lieu tied the letter to a proposed AI kill-switch bill .

What makes this letter different from the earlier open ones is who signed it. Not outside critics, but the engineers and executives closest to the frontier. One OpenAI safety researcher, Leo Gao, put the stakes bluntly: "the world is locked in a deadly race towards an intelligence explosion." His fix is the one the whole letter turns on. "To survive, we must coordinate to slow down the race."

Get the TNW newsletter

Get the most important tech news in your inbox each week.

Email address

Subscribe

Website

Thanks! Check your inbox to confirm your subscription.

By subscribing you agree to our Privacy Statement .

Get the TNW newsletter

✕ Close Get the TNW newsletter

Get the most important tech news in your inbox each week.

Email address

Subscribe

Website

Thanks! Check your inbox to confirm your subscription.

By subscribing you agree to our Privacy Statement .