The Off Switch

Why the world's most powerful AI is quietly turning from a normal product into something the government can switch off — and what that means for the rest of us.

By Ashok Murthy · IPNS Inc.

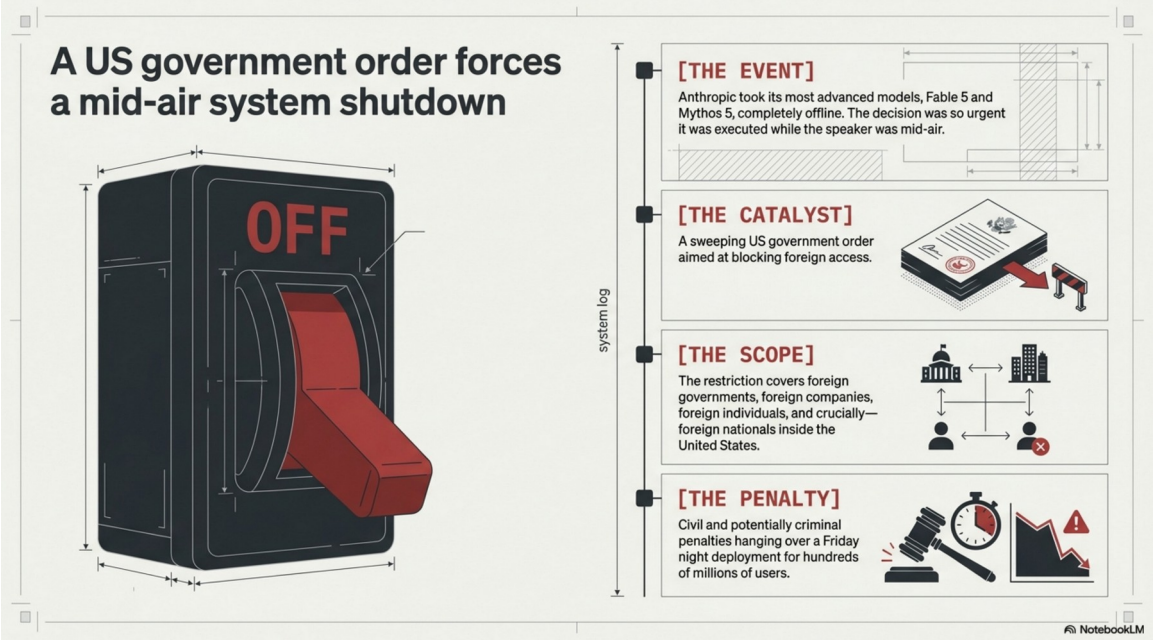

The scene: the smartest AI in the world, shut off in the middle of a sentence on a Friday night.

The story

Imagine a Friday night. The smartest AI in the world is in the middle of typing an answer for someone. Suddenly it stops. Not slowed down. Not "being retired next year." Just off. The reason isn't a crash. It's an order from the U.S. government, and the company pulls the plug right away — because if it doesn't, people could face huge fines or even criminal charges. The order sounds small: no foreign governments, no foreign companies, and — the key part — no "foreign nationals inside the United States" allowed to use it.

Here's the twist: I made that story up. But I didn't make up any of the parts that make it possible. **Every piece of it already exists in U.S. law. One version has already been switched on once. And right now, a real AI company is fighting something like this in court.** That's what makes it worth your attention.

That "foreign nationals" line is real law

That phrase sounds like a careful, targeted rule. It's actually one of the oldest and bluntest tools in trade law, and it has a name: a "deemed export." In plain terms — if you show certain controlled technology to someone who isn't a U.S. citizen, even if they're sitting in your own office, even if nothing ever leaves the building, the law treats it as if you shipped that technology to their home country.

This rule was written decades ago, for things like weapons blueprints. The new question experts are now asking: what if an AI's answers count as that kind of controlled technology? If they do, a non-citizen engineer who simply uses the model could be "exporting" something illegal without anyone meaning to.

Now think about how AI actually works. It's sold all over the world. The people who build it come from all over the world. It runs on cloud servers scattered across the globe and is reached through millions of connection points. So if a rule says you must guarantee that not

one single foreign national ever touches the model, there's no clean way to draw that line through the real system. The only way to be 100% sure is to shut the whole thing down. A small-sounding rule, applied to a worldwide system, becomes a giant off switch. The "narrowness" is the trap, not the safety net.

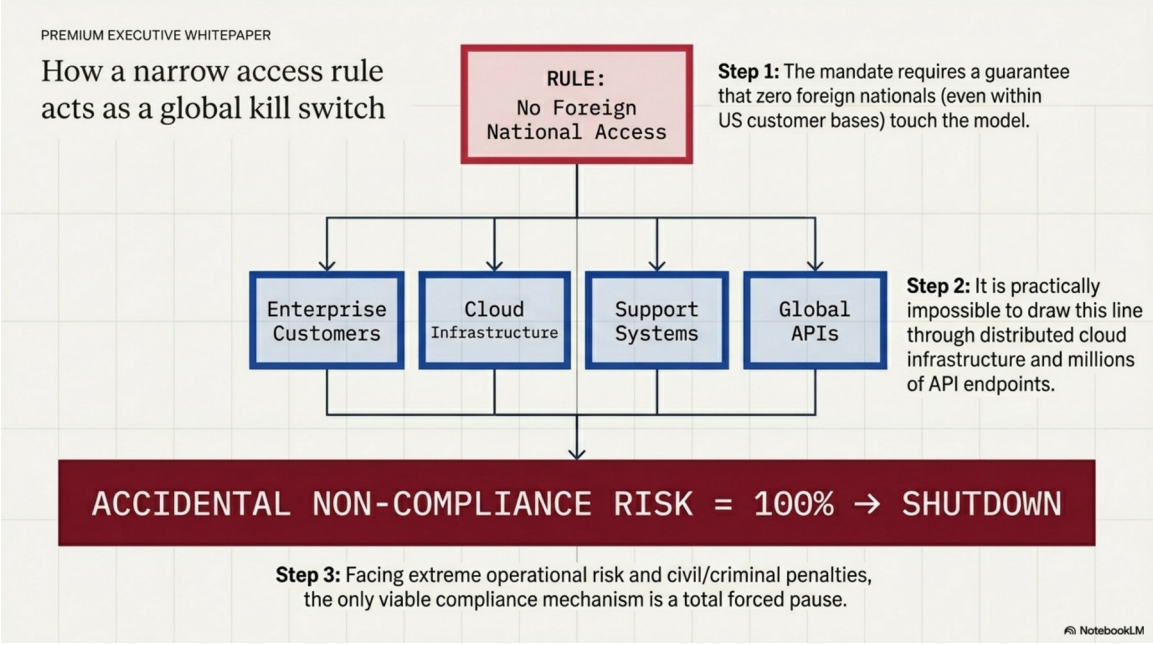

A small rule meets a worldwide system. The only way to fully obey it is to turn everything off.

It already happened once — then un-happened

If you think the rules here are steady, the last year and a half says otherwise. In January 2025, the government created a brand-new category that treated the "brains" of the most advanced AI models like a controlled export — something you'd need a license to send almost anywhere. For the first time, top AI was officially treated like a strategic national asset, not just software.

Then, two days before that rule kicked in, the government scrapped it. A replacement is still "coming." In the meantime, officials warned they'd come down hard on violations anyway. So the single biggest question for any company that depends on these models — is this a normal product I can use freely, or a tightly controlled asset? — got answered "controlled," then "free," then "we'll tell you later," all in one year. That's not solid ground to build a business on.

Can the government even do this? Maybe not so easily.

Here's the real danger. The government could freeze a private company's AI based on a vague warning — no clear law behind it, no published standard, no public review. Just power.

But there's good news, and it comes from a recent court fight. Any emergency "shut it down" order would likely lean on a 1977 law that's normally used for sanctions, called IEEPA. In February 2026, the Supreme Court looked at that same law and ruled, 6–3, that it does not give the president the power to slap on tariffs. The judges' reasoning: when the government wants to use broad, fuzzy wording to grab a huge new power, it needs Congress to have clearly handed over that power. If Congress didn't, the move fails. The Court even noted it was suspicious that no one had ever used the law that way in nearly 50 years.

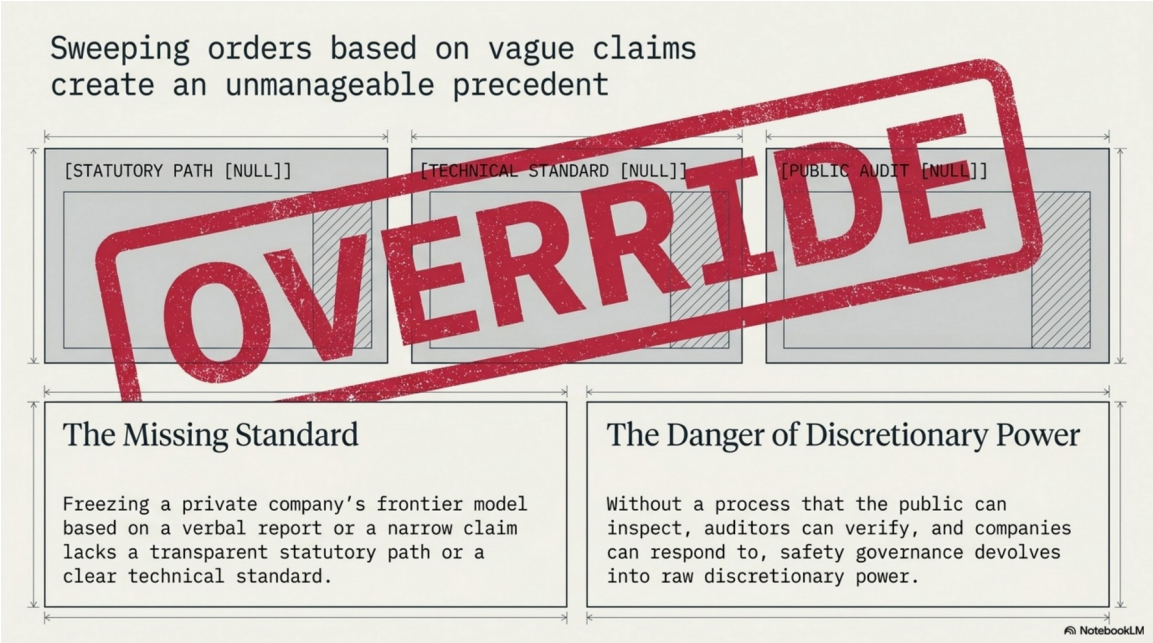

The worry: a freeze with no clear law behind it, no public standard, and no way to push back.

Apply that to AI. An order freezing a legal AI product on vague "national security" grounds, with no clear law behind it, is exactly the kind of sweeping move the Court just said it won't wave through. So that missing rulebook cuts both ways: it's a danger, but it's also the most likely reason an overreaching order would get struck down. The off switch is real — but if the government flips it without a fair, open process, a court can slap its hand.

This already happened to a real company

I've been saying "imagine." I shouldn't be. Through 2025, Anthropic — the maker of Claude — became the perfect example. It built special government versions of its AI, signed a two-year deal with the Defense Department worth up to $200 million, became the first to run a top-tier model on classified military networks, and got listed for easy purchase by federal agencies.

Then things turned fast. In early 2026, a disagreement over rules for using AI in weapons blew up. The government moved to label Anthropic a security risk and told agencies to stop using Claude. A few weeks later, in March 2026, a judge blocked that order and forced the government to put Anthropic back on its approved list. In just a few months, one company went from "trusted partner" to "security risk" to "put back by a judge." Betting everything on one company, in one shifting political mood, isn't a hypothetical risk. It already happened.

> *If your work depends on one model, one company, one country's mood, and one contract, you don't have a plan. You have a dependency.*

That word — dependency — is the whole point. For years, building on the best AI felt like plugging into electricity: always on, always there. The last year shows it's more like relying on a single supplier who can be cut off by a rule change, a contract fight, or a bad week in Washington. None of that is about whether the model is good. The model can be brilliant and still disappear from your work overnight. That's a new kind of risk, and it doesn't live in your code — it lives in law and politics, where most teams have never had to look.

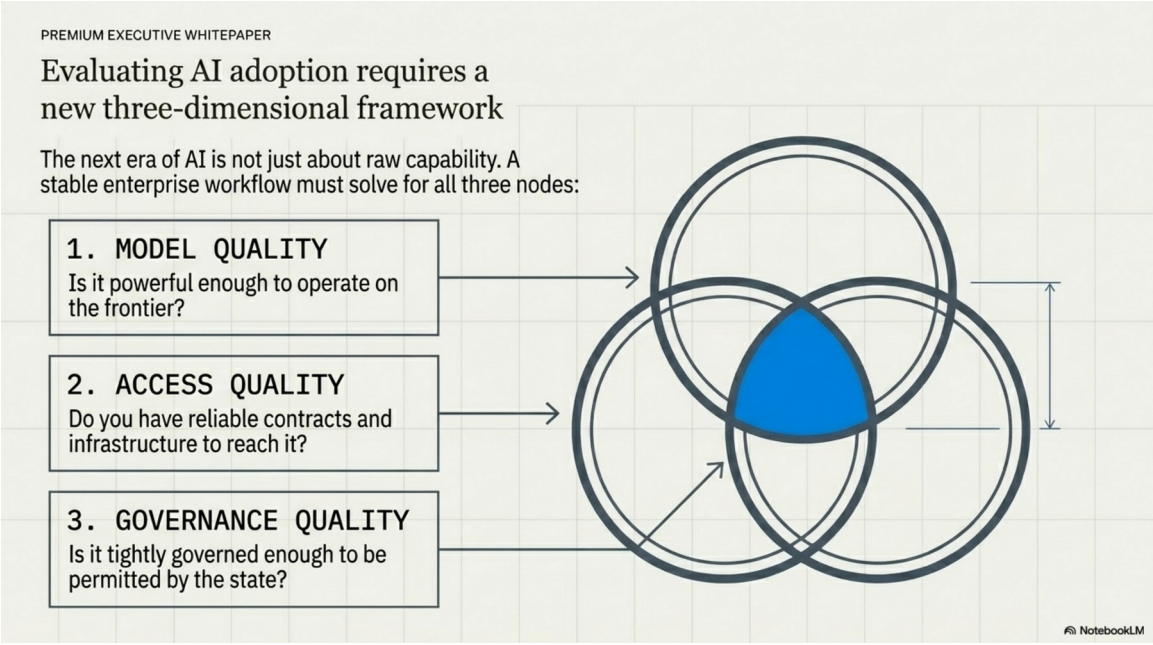

Being smart isn't enough anymore. A model also has to be easy to reach — and allowed to keep running.

What this means for you

If your work depends on AI — mine does — the lesson isn't political, and it isn't doom. It's practical. For years the only question was: is the model smart enough? Now there are three. Is it smart enough? Can you reliably reach it? And is it allowed to keep running? Chase only the first one, and you're one sudden rule change away from your tools going dark.

So do three things. First, know what you're leaning on — list every important task that depends on a single AI service, because that list is your risk. Second, keep a backup ready — set things up so your team can switch to another model in minutes, not months. Third, get your paperwork in order early — expect future rules to ask for proof of "trusted access," records of how the AI is used, and regular reports, and build that in now instead of scrambling later. The days of assuming the best AI will always be there, on yesterday's easy terms, are over.

The real question

It's tempting to call this "the end of unrestricted AI." That's close, but not quite. The shift — from treating top AI like a product to treating it like a controlled national asset — isn't on its way. It already happened: in the rules, in the contracts, and in the courtroom. The real question is simpler. When the government controls the off switch, will it use it through a fair, open process you can see and challenge? Or will it just be one phone call on a Friday night? If AI really is the engine of the future economy, then getting to it can't depend on one company staying on good terms with one administration. The smart move isn't only to ask for fair rules. It's to build your own setup assuming that off switch is real — and make sure your work survives the day someone reaches for it.

Sources

"Deemed export" rule — U.S. Export Administration Regulations, 15 C.F.R. Part 734.

Whether AI outputs count as controlled technology — Just Security and Lawfare, Dec. 2025.

The "brains/weights" export rule (ECCN 4E091), Jan. 15, 2025, and its reversal two days before taking effect, May 13, 2025 — U.S. Commerce Dept. (BIS); legal summaries by A&O Shearman, Finnegan, WilmerHale, Freshfields, Arnold & Porter.

The Supreme Court limiting emergency powers — Learning Resources / V.O.S. Selections v. Trump, Feb. 20, 2026 (6–3); background on IEEPA from CSIS; SCOTUSblog; Holland & Knight; Congressional Research Service.

Anthropic's government work, the security-risk label, and the March 26, 2026 court order restoring it — Anthropic; TechCrunch (Jun. 2025); Congressional Research Service (2026); Center for American Progress (Mar. 2026).