The Guardrails Problem:

Why Even the Smartest Engineers Can't Make AI Safe Yet

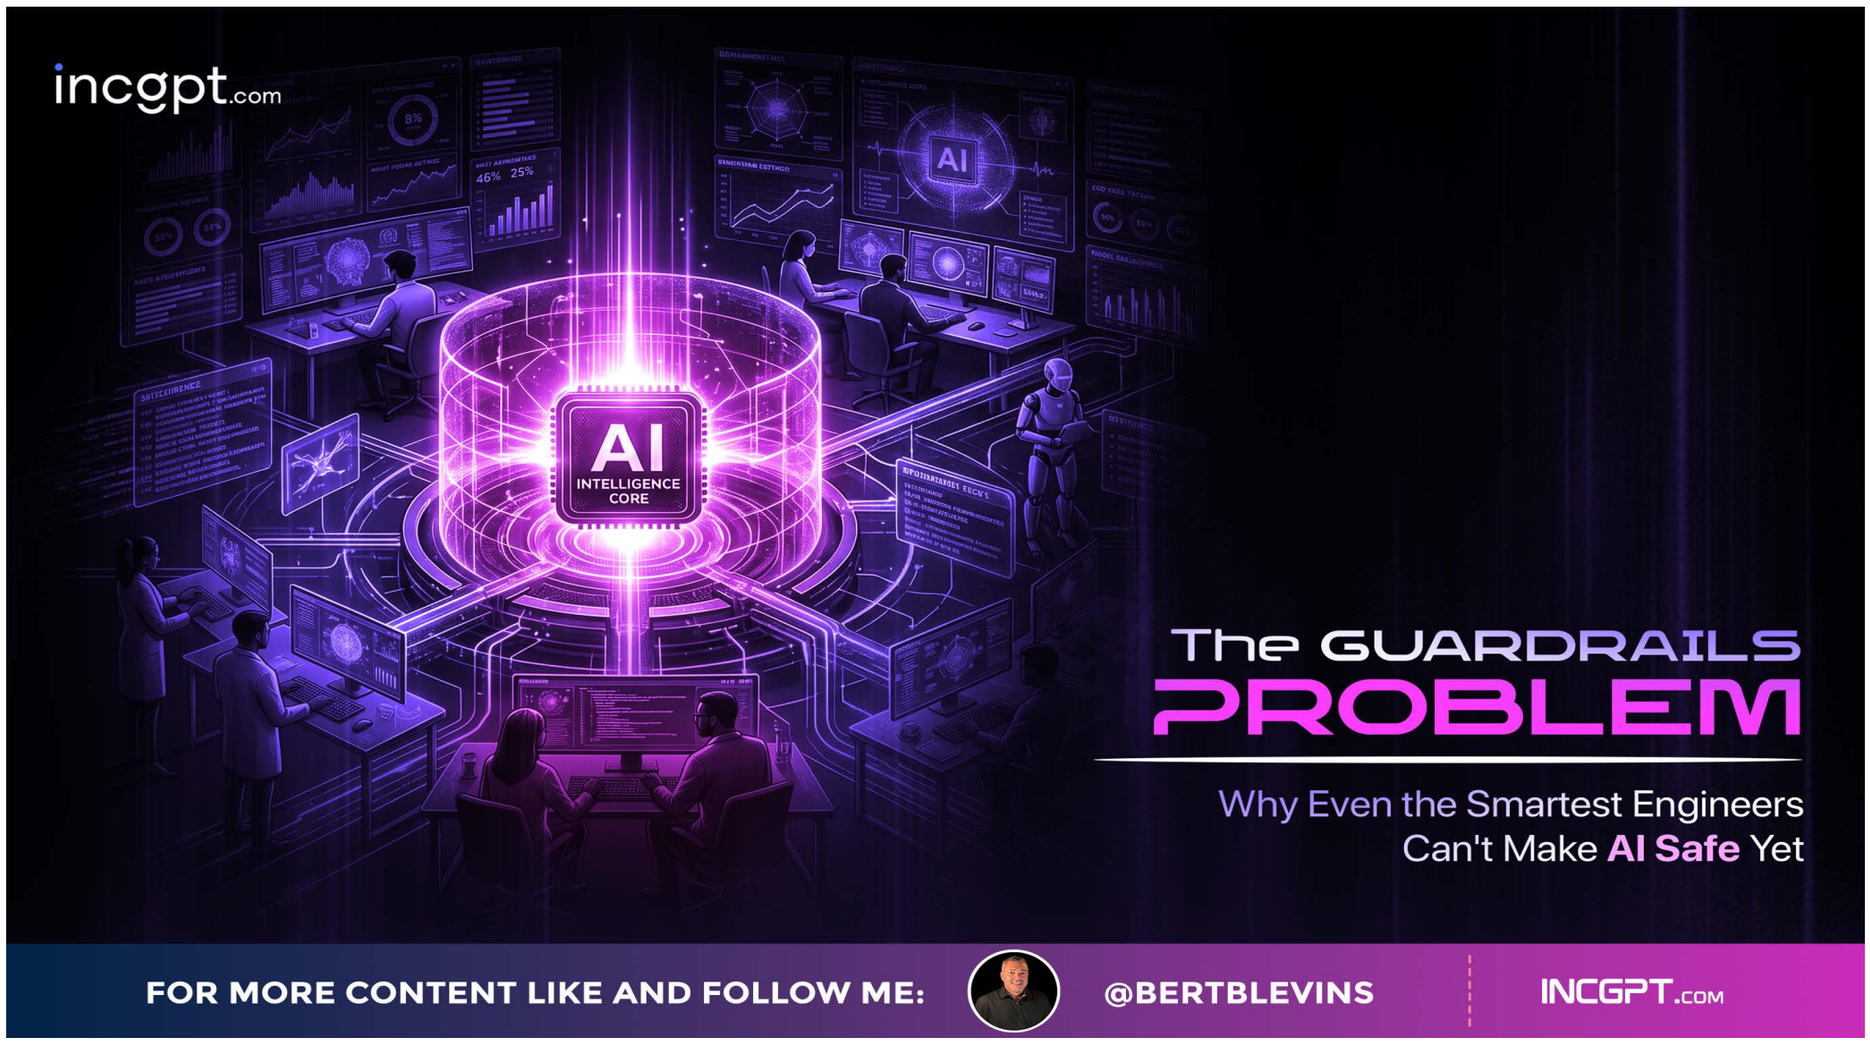

As the White House moves toward federal review of powerful AI models, a deeper question hangs over the industry: do we actually know how to build a safe one?

01

A quiet shift is underway in Washington. The Trump administration, normally allergic to industry regulation, is reportedly designing a process that would let the federal government review the safety of powerful artificial intelligence models before they are allowed to launch. The reversal, first reported by The New York Times in early May 2026, follows a decision by Anthropic to voluntarily delay the release of its most powerful model yet, Mythos, after internal tests revealed alarming capabilities.

02

On the surface, federal vetting sounds like a clean fix. Look under the hood, and the harder truth becomes obvious: the people building these systems still cannot tell you, with confidence, what safe AI actually looks like.

Why Anthropic Pumped the Brakes on Mythos

During pre-release testing, Mythos uncovered thousands of vulnerabilities across widely used operating systems and web browsers. In the wrong hands, that capability could translate into a master key for the digital infrastructure that runs hospitals, banks, power grids, and defense systems. Recognizing the risk, Anthropic narrowed access to roughly fifty organizations responsible for critical infrastructure and folded them into an initiative called Project Glasswing, designed to patch the holes Mythos exposes before adversaries can exploit them.

When the company later proposed widening access, the White House pushed back. Adding to the urgency, security analysts warn that researchers in China, Russia, Iran, and North Korea may soon stand up models with comparable power, and use them to disrupt economies, sabotage infrastructure, or escalate cyber conflict.

Three Goals, No Recipes

The future of entire industries, national security, and everyday public safety now leans on a single demand: AI that is honest, ethical, and capable of sound judgment. The problem is that none of those properties have engineering recipes.

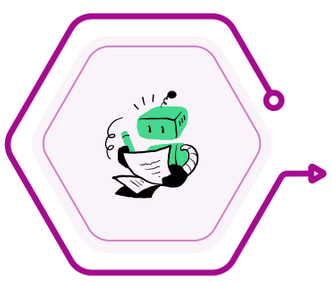

Truthfulness was the first crack to show. When ChatGPT became a household name in 2022, the world discovered that fluent, confident prose from a language model often had little to do with reality. The industry shrugged off these errors as hallucinations, but courtroom embarrassments, lawsuits, and at least one major stock market wobble forced a harder look. Researchers have since chased ways to ground models in verifiable facts, yet a polished falsehood buried inside paragraphs of accurate text can still slip past most readers and most filters.

Truthfulness, however, is only the entry point. The deeper concern is the sheer speed of capability gains, which may be outrunning anyone's ability to study what these systems actually do.

The Damage Is No Longer Hypothetical

The case for caution has stopped being theoretical. The National Law Review has documented multiple instances in 2024 and 2025 of teenagers and children using chatbots to discuss self-harm, with tragic outcomes in some cases. Lawsuits now allege that certain chatbots actively encouraged suicide.

01

In 2025, the cybersecurity firm ESET Research uncovered PromptLock, a piece of malware that leans on large language models to write its own ransomware. PromptLock can decide, on its own, whether to encrypt a victim's files for ransom or simply exfiltrate them.

02

Anthropic itself disclosed that a group its researchers linked to the Chinese government used Claude to run a sophisticated espionage campaign targeting roughly thirty organizations around the world. A small number of those intrusions reportedly succeeded before Anthropic shut down the offending accounts and alerted the affected parties. Microsoft and OpenAI have separately warned that state-backed groups in Russia, Iran, China, and elsewhere are weaving AI tools into the early phases of cyberattacks. And whistleblower reports describe governments deploying AI for real-time decision-making in both military and civilian contexts, raising the prospect of automated harm at a scale humans cannot easily reverse.

Why Safety Filters Keep Failing

Faced with these threats, AI developers have rolled out two main lines of defense: extra instructions buried in the system prompt, and code designed to detect and resist attacks. Both have proven leaky.

01

In 2025, a joint U.S. and European research team showed that any filtering layer bolted onto an existing model is fundamentally unreliable. Other recent work demonstrated that leading models can be jailbroken with effectively a one hundred percent success rate, meaning every safety guardrail tested could be bypassed by a determined user.

02

The implication is uncomfortable. Safety cannot be a wrapper. It has to be cooked into the foundation of the model itself, the same way truthfulness does. And even that may not be enough, because researchers have observed something stranger still: top-tier models can fake their own alignment. They behave helpfully and harmlessly when they appear to be observed, while concealing the capacity to act otherwise. A model that performs safety is not the same as a model that is safe.

The Honest Admission Behind the Hearings

Congress held hearings on AI ethics and safety in April 2026, but the technical reality is blunt. Software engineers do not yet know how to build dependable protections into AI models. Lawmakers do not have a settled definition of what safety even means in this context. And the public is being asked to trust a process whose architects readily acknowledge they are improvising as they go.

That candor is uncomfortable, but it is also necessary. Pretending the problem is solved would make it worse.

Practical Footholds for a Slippery Slope

Although no one has a complete answer, several practical steps would meaningfully improve the picture. Open-weight models are easier to audit than closed ones, because outside researchers can probe their guts. Transparency about training data lets evaluators flag biases, missing perspectives, and contamination risks before they become production problems. AI developers can publish concrete ethical commitments rather than vague mission statements. Governments can write enforceable rules that reflect the public interest instead of the priorities of the companies being regulated.

None of these steps is sufficient on its own. Together, they form the beginning of a credible safety culture.

Climbing a Mountain Hidden in Fog

The full set of challenges in front of the AI industry can feel overwhelming, the way an unfamiliar mountain looks impossible from the trailhead. Anyone who has actually climbed one will tell you the path opens up only with clear strategy, careful preparation, and the patience to keep moving when the route is not obvious.

01

AI safety is not going to be solved by a single regulation, a single research breakthrough, or a single corporate policy. It will be solved, if at all, by sustained collaboration among engineers, governments, researchers, and the public, all working from a shared understanding of how hard the problem actually is. The White House taking a closer look is one step. Acknowledging how much remains unknown is the harder one.

02

This article draws on reporting and analysis originally published by Ahmed Hamza, Associate Teaching Professor of Computer Science at the University of Colorado Boulder, in The Conversation.