The United States Constitution

Article II, Section 1 of the U.S. Constitution imposes only three eligibility requirements on persons serving as president, based on the officeholder's age, time of residency in the U.S., and citizenship status:

U.S. Constitution – Presidential Candidate Eligibility

"No person except a natural born Citizen, or a Citizen of the United States, at the time of the Adoption of this Constitution, shall be eligible to the Office of President; neither shall any person be eligible to that Office who shall not have attained to the Age of thirty-five Years, and been fourteen Years a Resident within the United States."

CORDOVA 2028 · POLICY BRIEF

Why Capability Denial Fails

An AI Security Doctrine Built on Collaboration, Not Containment

Completed publication · June 14, 2026

Core doctrine: America should not answer foreign advancement by restricting American capability. America should answer it by raising American readiness.

The Argument in One Paragraph

In June 2026, Anthropic said the United States government issued an export-control directive requiring the company to suspend access to Claude Fable 5 and Claude Mythos 5 for foreign nationals, including foreign nationals inside the United States. The company then disabled access broadly while it worked to comply. That action drew the security line in the wrong place. It restricted who could access a tool and where they were located, rather than focusing on the specific capability, the specific use, and whether the limit could actually be enforced. Capability denial - trying to make a country safer by giving its own builders, workers, researchers, students, and defenders less ability - fails for software knowledge that can be copied, rebuilt, and diffused. If America gates Americans while other nations continue advancing, denial becomes self-harm: the capability still spreads, but Americans fall behind. The durable path is broad access to general-purpose AI capability, universal literacy, hardened defenses, consistent capability-tiering, and narrow limits placed only where a genuine physical chokepoint makes them enforceable. AI is already being used to find software vulnerabilities. The question is not whether we can stop that. We cannot. The question is whether our defenders stay ahead of those who would do harm - and denial slows the defenders we can see while doing little to the actors we cannot.

What Happened

On June 9, 2026, Anthropic announced Claude Fable 5 and Claude Mythos 5. Anthropic described Fable 5 as a Mythos-class model for ambitious knowledge work and coding, with safeguards for cybersecurity and biology. It described Mythos 5 as its most capable model for cybersecurity and biology research, available only through trusted-access programs because of its dual-use risk.

Within days, Anthropic said the federal government issued an export-control directive requiring it to suspend access to Fable 5 and Mythos 5 for foreign nationals. Public reporting stated that the directive reached foreign persons inside the United States as well as abroad, including foreign-national employees of the company. Because the company could not immediately separate every permitted user from every restricted user in a compliant way, the practical result was that access was disabled broadly.

The reported trigger was a claim that Fable 5 safeguards could be bypassed for software-vulnerability discovery. The capability at the center of the dispute was having an AI read code and identify flaws in it - the same broad class of work defensive security researchers perform every day. Anthropic disputed that a narrow potential jailbreak justified

recalling a commercial model and pointed to comparable bug-finding capability in other AI systems. The company also continued to maintain that the highest-risk Mythos capability was already limited through trusted access.

Why this matters beyond one company: the precedent is the product, not the press cycle. A framework that restricts a diffusible capability at one domestic builder, while leaving comparable capability available elsewhere, becomes a template for how the United States governs AI. That template deserves scrutiny on its own terms, independent of any one firm's commercial interests.

The Structural Flaw: Denial Fails for Anything That Diffuses

A capability you can restrict is one you can contain. But the capabilities in question are not physical objects. They are knowledge of how to do something, expressed in software. Knowledge of that kind diffuses. It is published, replicated, leaked, distilled, fine-tuned, and rebuilt. The marginal cost of copying it approaches zero. When a government restricts a diffusible capability at one source, it does not remove the capability from the world. It removes it from the actor it can see and regulate, while actors it cannot regulate proceed unaffected.

This is not a new lesson. In the 1990s, the United States treated strong encryption as a controlled export. The underlying mathematics was published and rebuilt regardless, and the rules were later relaxed for publicly available and mass-market encryption. The pattern is the same here: a control that can be routed around buys little durable safety and imposes real cost on the people who obey it.

The Distinction That Makes the Doctrine Credible

Denial is not always futile - and an honest doctrine must say where it works. The decisive variable is whether a capability is gated by diffusible knowledge or by a physical chokepoint.

Diffusible capabilities - code, algorithms, model techniques, prompts, workflows, and knowledge - cannot be durably contained. Restriction here mostly penalizes the compliant.

Chokepoint-gated capabilities - those that require scarce physical inputs, specialized equipment, controlled materials, or a very small number of production facilities - can sometimes be meaningfully constrained, because there is a real bottleneck to hold.

A serious policy distinguishes the two and is honest about it. A doctrine that claims denial never works collapses the moment someone points to chip controls, fabrication equipment, or controlled physical inputs. A doctrine that says denial fails for software and knowledge specifically, while conceding that physical chokepoints can sometimes be held, stands on defensible ground. The June 2026 order failed this test: it tried to fence a diffusible capability - reading code to find flaws - as if it were a physical good.

Selective Application Based on the Public Record

Based on the public record so far, the line was drawn unevenly. Anthropic was restricted after a concern over software-vulnerability discovery, while public reporting and technical discussion described comparable bug-finding capability in other AI systems. If the capability is the danger, then restricting it at one builder and not across equivalent systems does not reduce the danger; it relocates the burden to the compliant builder and its users.

The better instinct is to gate the riskiest capability tier itself, not the nationality of a user or the location of a company. Anthropic's own trusted-access approach for Mythos 5 points in that direction: the highest-risk capability is limited to vetted security and research partners. That model still needs public standards, due process, and accountability, but its organizing question is the right one: what can this capability do, and under what conditions should it be used? A

nationality-based shutdown asks a different question. It asks who may touch the tool. That is a weaker security standard.

The Real Question: Do Defenders Stay Ahead of Attackers?

AI is already being used to find vulnerabilities in software. That is not a future risk to be prevented; it is a present reality to be managed. Framing the challenge as "are we ready" treats safety as a finish line. It is not. Safety in this domain is a race condition - a continuous contest between defensive use and offensive use - and the policy question is which side our choices accelerate.

Restriction loses that race by construction. It slows the defenders who operate in the open and submit to oversight, while leaving untouched the actors who never asked permission. The same capability that can probe a system for weaknesses is the capability that lets defenders find and close those weaknesses first. During Project Glasswing, Anthropic reported that roughly 50 partners used Mythos Preview to identify more than 10,000 high- or critical-severity vulnerabilities across important software, with several partners reporting hundreds of high-severity findings. Those numbers are not an argument for carelessness. They are an argument for scaling defense faster than offense.

The National Disadvantage Risk

The danger is not only that bad actors gain capability. The danger is that Americans are denied capability while the rest of the world continues moving. A policy that gates American builders, workers, researchers, students, small businesses, and defenders while foreign competitors advance without the same limits does not create safety. It creates national disadvantage.

In a new technological era, the United States cannot protect its people by making them less capable. Our duty is to advance Americans at the same level as, or ahead of, those who would use these systems against us - with literacy, accountability, defensive strength, and narrow controls where controls can actually be enforced. America should not answer foreign advancement by restricting American capability. America should answer it by raising American readiness.

The Alternative: Four Pillars of a Collaboration Doctrine

If denial fails for diffusible capability, the answer is not to pretend there is no risk. It is to build a posture that takes risk seriously without surrendering the benefit. Four pillars hold it up.

1. Tier by capability, not by identity. Most of what makes these tools transformative - drafting, tutoring, research, code review, logistics, medical administration, education, accessibility, and ordinary business productivity - carries no catastrophic risk and should be widely available. The genuinely dangerous slice is narrower and more specific. Restrict that slice at the capability tier for everyone, with transparent standards and vetting for the highest-risk uses, rather than using nationality or company identity as a substitute for technical judgment.

2. Universal risk literacy. Broad access should come with broad understanding. The public should be trained to recognize accident-shaped dangers: data exposure, over-trust in outputs, mishandling sensitive results, false certainty, unsafe automation, and weak review practices. Literacy is a real mitigation for the well-intentioned, but it must carry only the weight it can bear. It does not, by itself, stop a determined bad actor. It must be paired with hardened systems.

3. Defense and resilience investment. Because dangerous knowledge will diffuse regardless of any one restriction, the responsible assumption is that it gets out - and the work is to harden the targets. Public investment in defensive security, vulnerability remediation, secure-by-design software, resilient infrastructure, and rapid patching is what converts "the capability exists" from a crisis into a manageable condition. This is where a sovereign government's effort belongs: not in denying its own people a tool, but in making the systems they depend on harder to break.

4. Narrow, enforceable limits at real chokepoints. Where a genuine physical bottleneck exists, a limit can be held and may be worth holding. The discipline is to reserve hard restriction for those cases, to review them regularly, and to refuse the comforting illusion that fencing a diffusible capability at a domestic builder makes anyone safer.

A Governing Standard for Future AI Restrictions

No AI capability restriction should be imposed unless the government can satisfy a clear public standard, subject to narrow national-security redactions where truly necessary:

Identify the specific capability being restricted, not merely the company or model name.
Show why the capability creates a concrete risk beyond ordinary dual-use concern.
Apply the standard across equivalent systems, domestic and foreign, rather than singling out one builder.
Explain why the restriction is technically enforceable and not merely symbolic.
Provide a response process for the affected party and a path to restoration when safeguards are improved.
Use the narrowest restriction capable of addressing the specific risk.
This standard does not eliminate government authority. It disciplines it. It says the United States can act against real danger, but it must act in a way that strengthens American readiness instead of weakening it.

Change Is Here: AI as a Collaborator

The deepest divide in AI policy is between two theories of safety. One holds that the way to be safe is to have less capability - and so it denies, contains, and excludes. The other holds that the way to be safe is to have more capability, broadly distributed, with defensive use deliberately outpacing offensive use. The first theory produced the June 2026 shutdown, which helped no one with good intentions and left every actor with bad intentions exactly where they were. The second theory - collaboration - is the one that matches the technology actually in front of us.

We will have risk in anything new. That is the price of every tool worth having. The task is not to refuse the tool but to decide, together, how we carry the risk: who is trained, what is hardened, where the few real limits sit, how those limits are reviewed, and how the benefit reaches everyone with good intentions rather than only those who can afford enterprise access. A Cordova administration treats AI as a collaborator to be integrated with accountability - not a threat to be contained at the expense of our own people's capability.

What This Brief Claims, and What It Does Not

It claims: that the June 2026 order was applied unevenly based on the public record; that it restricted a diffusible software capability; that denial of diffusible capability fails as a durable security strategy; that limiting Americans while other nations continue advancing creates national disadvantage; and that capability tiering, public literacy, defensive resilience, and narrow chokepoint limits are the defensible alternative.

It does not claim: that any company was deliberately persecuted, that the order was driven by commercial or political motive, or that all AI restrictions are illegitimate. Those questions are contested or unproven, and this argument does not depend on them. The case stands on the public facts and the structure of the technology itself - which is precisely what makes it durable.

Sources and Factual Basis

Factual basis current through June 14, 2026. Source notes are provided for publication review and link verification.

1. Anthropic, “Claude Fable 5 and Claude Mythos 5,” June 9, 2026. https://www.anthropic.com/news/claude-fable-5-mythos-5

2. Anthropic, “Claude Fable 5” model page, including June 12 availability update. https://www.anthropic.com/claude/fable

3. Anthropic, “Claude Mythos 5” model page, including June 12 availability update and trusted-access description. https://www.anthropic.com/claude/mythos

4. Reuters, “Anthropic disables top-tier AI models after US order limiting foreign access,” June 13, 2026. https://www.reuters.com/technology/us-blocks-foreign-access-anthropics-most-advanced-ai-models-axios-reports-2026-06-13/

5. Reuters, “EU Commission looking at practical consequences of Anthropic decision, spokesperson says,” June 14, 2026. https://www.reuters.com/legal/litigation/eu-commission-looking-practical-consequences-anthropic-decision-spokesperson-2026-06-14/

6. Anthropic, “Project Glasswing: An initial update,” May 22, 2026. https://www.anthropic.com/research/glasswing-initial-update

7. Anthropic, “Expanding Project Glasswing,” June 2, 2026. https://www.anthropic.com/news/expanding-project-glasswing

8. Tom’s Hardware, “U.S. gov’t orders Anthropic to disable its newest AI models worldwide due to security threats,” June 13, 2026. https://www.tomshardware.com/tech-industry/artificial-intelligence/us-export-control-order-forces-anthropic-to-disable-claude-fable-5-and-mythos-5-worldwide

9. Bureau of Industry and Security, “Encryption items not subject to the EAR,” public guidance. https://www.bis.gov/learn-support/encryption-controls/encryption-items-not-subject-to-ear

10. Electronic Frontier Foundation, “U.S. Export Controls and ‘Published’ Encryption Source Code, Explained,” Aug. 27, 2019. https://www.eff.org/deeplinks/2019/08/us-export-controls-and-published-encryption-source-code-explained

Source caveat: capability-equivalence between models is a contested public claim. This brief treats it as a public-record concern sufficient to test whether a restriction is being applied consistently, not as an independently adjudicated technical