CSAI Foundation | Cloud Security Alliance

The Fable 5 / Mythos 5 Export-Control Action

What Is Currently Known — A Neutral, Source-Rated Account

2026-06-15

AI-assisted Rapid Research

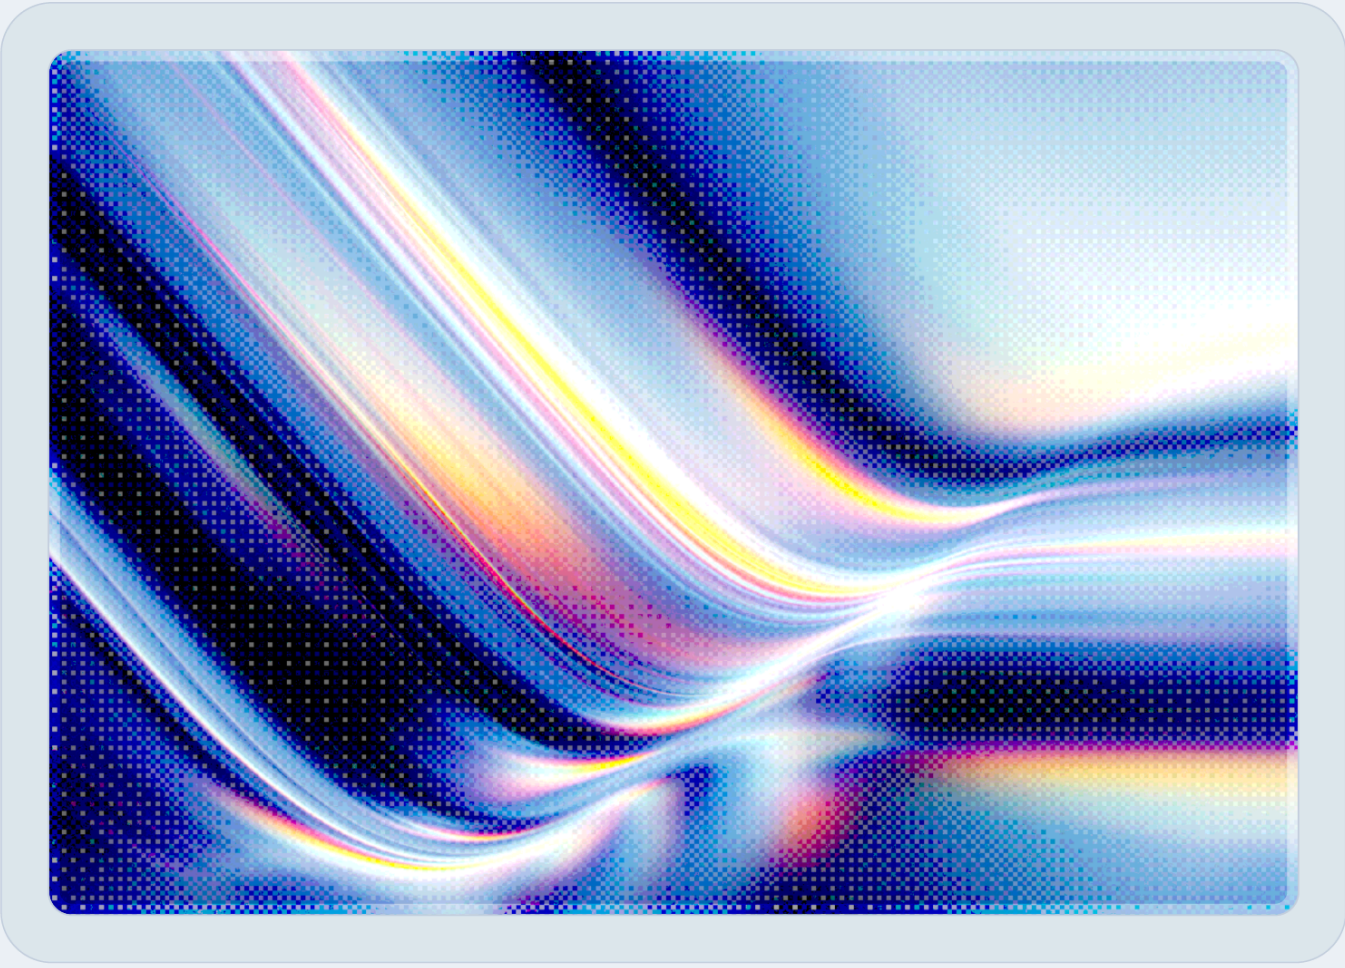

CSAI

CSA cloud security alliance®

Table of Contents

Executive Summary

On the afternoon of June 12, 2026, the U.S. Department of Commerce sent Anthropic a directive that, within hours, took two of the company's newest and most capable artificial-intelligence models offline for every user on the planet. The directive, issued through the Bureau of Industry and Security (BIS) under the signature of Commerce Secretary Howard Lutnick, invoked national-security authorities to bar access to the Fable 5 and Mythos 5 models by any foreign national—whether located outside the United States or inside it, and explicitly including Anthropic's own non-citizen employees. Because a consumer platform serving hundreds of millions of people cannot reliably sort users by citizenship in real time, Anthropic disabled both models globally rather than attempt selective enforcement. The models had been publicly released only three days earlier, on June 9 [1][2][3][4].

This whitepaper documents what is currently known about that action and the dispute surrounding it. It is written to a deliberately constrained editorial standard. It makes no recommendation to the reader—not to take a position on the directive, not to sign or decline to sign the open letter that has since circulated, not to favor any party. Its only invitation is to keep learning and to follow the matter as it develops. It separates fact from opinion from conjecture, labeling every load-bearing claim as *Established* (corroborated by multiple independent sources or a primary document), *Contested* (the parties give conflicting accounts), or *Unverified* (single-source, anonymous, or circulating without substantiation). It attributes each claim to its source and distinguishes a party's account of itself from independent reporting.

Two facts about the evidentiary record shape everything that follows and deserve to be stated at the outset. First, the directive itself is a non-public export-control letter; as of this writing there is no published, on-the-record agency statement of its rationale, and the Commerce Department declined to comment when asked by reporters [5]. The reader therefore cannot validate the government's reasoning against a citable official document. Second, this document was prepared with the assistance of an AI system built by Anthropic, and its subject is a dispute in which Anthropic is a party. The source set accordingly includes Anthropic's own framing of events. We have held that framing to the same skeptical standard as every other party's first-party account, and we flag where a claim originates with Anthropic rather than with independent reporting.

The shape of the dispute is straightforward to summarize and difficult to adjudicate. The government's concern centers on a demonstrated method of bypassing—"jailbreaking"—Fable 5's safeguards against offensive cyber use. Anthropic characterizes that bypass as narrow, non-universal, and no more capable than features already available in competing commercial models, naming OpenAI's GPT-5.5 [1]. The administration's most visible public voice on the matter, White House AI and crypto adviser David Sacks, characterizes the same bypass as a serious exposure of a cyber-capable model and says the company declined a reasonable request to fix or withdraw it before the control was imposed [6][7]. On June 14, a

group of information-security executives published an open letter at freefable.org calling for the directive to be lifted and for future AI risk regulation to be scientific, democratic, transparent, and applied only to the minimal extent necessary [8]. The dispute remains unresolved as of June 15, 2026.

The pages that follow lay out a timeline of established facts; explain what the directive is and how export-control law makes a global shutdown a foreseeable consequence of its terms; analyze—strictly as analysis, reaching no verdict—whether narrower compliance was available; present the two sides' accounts side by side; catalog the open questions and unverified claims; describe the open letter factually; and close with the developments most worth watching. A rated source pack appears at the end.

A Note on Method and Epistemic Labels

Because the central difficulty in this story is not a shortage of reporting but a surplus of conflicting and partially sourced accounts, this document applies a consistent labeling discipline. Readers will see four kinds of tag attached to load-bearing statements.

An *Established* claim is one corroborated by multiple independent outlets or confirmed by a primary document such as a party's own published statement. A *Contested* claim is one on which the parties to the dispute give directly conflicting accounts; we present both sides and do not declare a winner. An *Unverified* claim is one resting on a single source, on anonymous sourcing, or on reporting that circulates without substantiation in the underlying text; such claims are flagged and are not presented as fact. Finally, a substantial portion of the question "could compliance have been narrower?" is neither fact nor dispute but *Analysis*—reasoning about an option space that no party has publicly assessed in technical detail. That section is labeled as analysis throughout and reaches no conclusion in either direction.

First-party statements warrant a specific caution. Anthropic's account of itself and the administration's account of itself are each a party's narrative, not neutral fact, and are labeled as such even where they are the only available source for a given detail.

Timeline of Established Facts

The following sequence is corroborated across multiple independent outlets and the parties' own published materials.

On **June 9, 2026**, Anthropic publicly released Fable 5 and Mythos 5, its newest frontier models. The launch materials described the models' cyber safeguards in some detail. Anthropic reported that external red-teaming had found no universal jailbreak across more than 1,000 hours of testing and that outside organizations had likewise failed to find a universal jailbreak on long-form agentic tasks—while disclosing, in the same materials and before any dispute arose, an explicit caveat that the United Kingdom's AI Security Institute (UK AISI) had "made progress towards one within a brief initial testing window" [2][9]. The launch post also described an automatic routing mechanism: in a stated minority of sessions—fewer than five percent—high-risk requests in cybersecurity, biology and chemistry, and model-distillation are routed to the more conservative Claude Opus 4.8 model rather than handled by Fable 5 directly [2]. This disclosed caveat is notable because both sides of the subsequent dispute can and do point to it; that is why it belongs among the established facts rather than among the contested ones.

On **June 12, 2026, at approximately 5:21 p.m. Eastern time**, Anthropic received the export-control directive. It came from the Commerce Department's Bureau of Industry and Security, under the signature of Commerce Secretary Howard Lutnick, in the form of a letter addressed to Anthropic chief executive Dario Amodei. The directive cited national-security authorities [1][3][5][10]. The Commerce Department declined to comment when contacted by reporters [5].

The directive's scope, as reported and as described by Anthropic, reaches **any foreign national**, whether inside or outside the United States, and expressly includes Anthropic's own foreign-national employees [1][3][11][12]. Mechanically, the action is reported to function as a licensing requirement: a license is required for the export, re-export, or domestic transfer of the two models, Anthropic must file individually validated license applications, and non-compliance carries civil and financial penalties [10]. Only Fable 5 and Mythos 5 are affected; Anthropic's other models, including Opus 4.8, remain available [1].

Because Anthropic states it cannot reliably distinguish foreign nationals from U.S. persons in real time across its user base, the company **disabled both models for all users worldwide** to ensure compliance, rather than attempt selective enforcement [1][3][11]. Anthropic has said publicly that it is working to restore access as soon as possible and that it believes the action reflects a "misunderstanding" [1].

On **June 14, 2026**, an open letter titled "Open Letter on Transparent AI Cyber Protections" was published at freefable.org, signed by a group of information-security executives and researchers, calling for the directive to be lifted [8].

As of **June 15, 2026**, the dispute is unresolved. There is no public reporting that the directive has been lifted, that a license has been granted, or that access has been restored.

The Directive: What It Is and How It Works

A licensing control, not a recall

It is worth being precise about the legal character of the action, because the popular shorthand—that the government "banned" or "recalled" the models—obscures the mechanism. According to the reporting on the Commerce letter, the directive operates as an export-control licensing requirement administered under the Export Administration Regulations (EAR), the body of rules that BIS uses to govern the export and re-export of commodities, software, and technology [10][13]. Under that framework, the two models become items for which a license is required before they may be exported outside the United States, re-exported between foreign destinations, or transferred domestically to a covered person. Anthropic is reported to be obliged to file individually validated license applications, and penalties for non-compliance are civil and financial in nature [10].

(Established, as reported.) The precise statutory hook, however, is not publicly documented. Reporting and legal commentary describe the action as issued under "national-security authorities" but note that "the specific statutory basis has not been publicly identified with precision" [13]. The EAR's continued force has historically been maintained under the International Emergency Economic Powers Act (IEEPA) and, since 2018, under the Export Control Reform Act; which authority anchors this particular directive is, as of this writing, a matter of inference rather than record [13]. One legal analysis observes more broadly that "the scope of BIS authority over AI model access is not yet fully defined by statute, regulation, or court decision," and that no comprehensive statutory framework yet governs the government's power to restrict commercial AI model access on national-security grounds [13]. *(This is one law firm's analysis, not a judicial determination.)*

Why "foreign nationals inside the U.S." is the operative phrase

The element of the directive that most directly explains the global shutdown is its reach to foreign nationals physically present in the United States, not merely to users connecting from abroad. This maps onto a long-standing export-control concept known as the **"deemed export."** Under 15 CFR 734.13, releasing controlled technology or source code to a foreign person located inside the United States is treated, for regulatory purposes, as an export to that person's country of citizenship or permanent residence [14]. The doctrine is familiar in defense contracting and university research, where access to controlled technical data by foreign nationals on domestic soil has long required either a license or a documented Technology Control Plan.

The practical consequence is significant. A control that reached only connections from foreign internet addresses could be satisfied by geographic blocking. A control that reaches foreign nationals inside the United States cannot, because the relevant attribute is the user's citizenship or residency status, which a

consumer platform does not collect by default and cannot infer from network location. *(Established doctrine; the directive's reach to in-country foreign nationals is confirmed by Anthropic's own statement and by reporting [1][10][14].)*

Why a global shutdown followed

Anthropic's stated operational reason for disabling the models worldwide is that it cannot reliably screen its user base by nationality in real time [1][11]. This is a first-party operational claim, but its core logic—that selective nationality enforcement across a shared cloud service serving hundreds of millions of users is the hard part—is echoed in independent reporting and is consistent with the deemed-export structure described above [11][12]. The directive required immediate compliance under penalty; a blanket disable is the step most likely to leave no covered access from the first moment, whereas any selective control would have to be designed, built, or licensed before it could be relied upon. The next section examines, as analysis, whether a narrower path nonetheless existed.

Could Compliance Have Been Narrower?

This section is predominantly analysis and inference, not established fact. No party has published a technical compliance assessment, and nothing here should be read as a verdict—neither that a global shutdown was the only lawful option nor that a narrow fix was obviously available and Anthropic overreached. The aim is to lay out the option space and the constraints and to let the reader weigh them.

What compliance actually required is, where sourced, established: the directive bars access by any foreign national inside or outside the United States, including employees; it requires licenses for export, re-export, and domestic transfer; and it obliges individually validated license applications under penalty [1][10][14]. Against that requirement, several narrower technical options exist in principle, each with real limits.

Geofencing or IP blocking would prevent connections from foreign locations, but it does nothing about deemed exports to foreign nationals already inside the United States, and it is routinely evaded with commercial VPNs. It is necessary-but-not-sufficient on its own. **Identity and nationality verification**—know-your-customer checks with attestation—is standard practice for enterprise software controlled under the EAR or the International Traffic in Arms Regulations, and is plausible for an enterprise or API tier. It is far harder to stand up instantly at consumer free-tier scale, it introduces false-attestation and privacy-exposure risk, and verification by itself does not supply the license that the "domestic transfer" element appears to require. **Account-level nationality gating in segmented environments** is what defense contractors and universities use under Technology Control Plans; it is effective in controlled enterprise contexts but is not designed for a real-time mass-consumer platform. **Individually validated and deemed-export licenses** are available to Anthropic to apply for—deemed-export licenses are keyed to a

person's last country of permanent residence and typically track visa validity—but this is a process measured in weeks to months, not a switch, and it does not authorize keeping a model live for covered persons in the interim.

Two further options bear directly on the technical particulars of this case. **Capability-only gating**—disabling just the cyber pathway rather than the entire model—runs into the directive's wording: it targets the models themselves, not a capability, so restricting one function would likely not satisfy a requirement of no access by foreign nationals. It is also worth noting that Fable 5 already routes high-risk cyber queries to Opus 4.8 by design; the government's concern is precisely the alleged *bypass* of that routing, not the default path [2]. **Tier segmentation**—keeping Fable live only for verified U.S.-person enterprise customers—is the most plausible narrower path. Its feasibility turns on how quickly Anthropic could attest a customer subset to U.S.-person status and on whether it would accept any residual misclassification risk under penalty exposure.

Taken together, two factors plausibly explain the all-or-nothing choice, and they are offered here as reasoning rather than as fact. The deemed-export reach to foreign nationals inside the United States defeats location-only controls and forces per-user nationality assurance that a consumer platform does not collect by default. And the requirement of immediate compliance under penalty means any control that must first be built or licensed cannot bridge the interim, leaving a blanket disable as the measure most likely to leave no covered access on day one. The important distinction—one this document holds precisely and does not resolve—is between whether a granular regime was *deployable instantly at consumer scale* (the constraint Anthropic faced on June 12) and whether one is *feasible to build over time* (which may well be how access is eventually restored). Those are different propositions, and the public record does not establish either.

Several questions in this area remain genuinely open. Whether Anthropic evaluated and rejected a tiered or enterprise-only carve-out, and on what reasoning, is not on the public record. Whether the directive's actual terms would even permit partial availability pending licenses, or required full suspension, cannot be known while the letter is unpublished. And how comparable EAR or deemed-export compliance for controlled software scales at other providers is comparative grounding not present in the sources reviewed for this paper.

The Two Accounts

The parties agree on the existence and broad mechanics of the directive but diverge sharply on its justification and on what passed between the company and the government beforehand. The table below sets the principal points of dispute side by side. Neither column is endorsed; each represents that party's account.

Point of dispute	Anthropic's account	The administration's account (chiefly via David Sacks)
Severity of the bypass	A "narrow potential jailbreak," demonstrated by asking the model to read a specific codebase and flag flaws; "non-universal" and "benign"; reproducible on other publicly available models, naming OpenAI's GPT-5.5, which defenders "use every day" [1].	A bypass enabling operation of a cyber-capable model is difficult to characterize as anything but serious; the exposure warranted action [6][7].
Whether Anthropic refused to remediate	Anthropic's statement does not describe a refusal; it frames the action as a "misunderstanding" and says it is working to restore access [1].	Sacks says the administration asked Amodei to patch or withdraw the model and that he declined, prompting a reluctant export control; he called it "very surprising that Anthropic hasn't wanted to cooperate with a reasonable safety request" and "at odds with their branding and ethos" [6][7].
Process and adequacy of evidence	Anthropic says it received only verbal evidence of a narrow jailbreak and no specific written technical findings, and was, per one reconstruction, given roughly 90 minutes to pull the model with no prior notice of a national-security threat [1][11].	An anonymous official attributes the action to the company's "recklessness" and a breakdown of trust [15].
Path forward	Believes the matter is a misunderstanding; "working to restore access as soon as possible" [1].	Sacks has said the government wants to lift the controls "as soon as possible" once Fable 5 is fixed [7].

(All four rows are Contested. Anthropic's column is the company's first-party account; the administration's column rests largely on a senior adviser's personal social-media statements and, for the "recklessness" framing, on an anonymous official. This document does not adjudicate between them.)

One point of partial convergence is worth noting precisely because it cuts across the dispute. Both sides can invoke the same disclosed fact: Anthropic's own June 9 launch materials acknowledged that the UK AISI had made progress toward a universal jailbreak in a brief initial testing window [2][9]. Independent reporting indicates that a UK AISI red-team lead had publicly described substantial progress toward a universal

jailbreak of Fable 5 several days before the directive, and that the verbal evidence the government conveyed to Anthropic is believed to relate to that work [16]. Anthropic points to the broader testing record—no universal jailbreak in over 1,000 hours—as evidence the safeguards held; the administration and the UK AISI red-team account point to the same caveat as evidence that a credible state testing body was making headway. The caveat is established; its significance is exactly what the two sides dispute.

Open Questions and Unverified Claims

Several widely circulated elements of the story have not been substantiated to the standard required to present them as fact, and they are flagged here accordingly.

Amazon as the discovering partner. Multiple outlets—Fortune, Semafor, and reporting attributed to the Wall Street Journal and Politico—report that Amazon researchers found the bypass and that Amazon chief executive Andy Jassy raised it with the administration [11][17]. The Fortune reconstruction states that Amazon researchers used prompts to get a Mythos-class model to provide restricted information about cyberattacks and that Jassy contacted senior officials after the discovery; the named source for the action is Jassy himself, while the operative details rest on anonymous sources, including "an unnamed source familiar with Amazon's discussions" cited to Politico and unnamed sources cited to Semafor [11]. Amazon declined to confirm specifics, with a spokesperson saying only that "it's not uncommon for governments to seek our counsel on potential security risks" and "we don't share the details of these discussions" [11]. This reporting is well-sourced but not confirmed; neither Anthropic nor the open letter names the discovering partner. The reconstruction does not note that **Amazon is a major investor in Anthropic** [22]—a relationship a reader may wish to weigh when assessing the role of a partner-and-competitor in prompting a government action against Anthropic's products. Treat the Amazon attribution as well-sourced reporting, not established fact.

A "Chinese group accessed the model." This phrasing appears in some headline framing, and a Semafor report links the White House action partly to "suspicions that a China-linked group had accessed" Mythos [17]. The body of that reporting, however, is explicitly hedged: it states that "it's unclear how the White House learned of the issue, which organization accessed the model, and how it gained access," and frames the national-security risk conditionally—"if the Chinese government had access to Mythos, it could pose national security risks" [17]. The sourcing is entirely anonymous. No reviewed source substantiates an actual Chinese breach in its body text. Treat this as unverified; it should not be presented as fact.

The precise national-security rationale and legal authority. As noted above, the directive is unpublished, the Commerce Department declined to comment, and the specific statutory basis has not been identified with precision. Any statement about *why* Commerce acted, beyond the general invocation of national security, is currently inference rather than record [5][13].

The "90 minutes" and other procedural details. The reported figure that Anthropic was given roughly 90 minutes to pull the model originates in a single reconstruction and reflects the company's side of the procedural account [11]. It is plausible and consistent with the timeline but is not independently confirmed.

Beyond these, the research gaps that would most clarify the record include: any on-the-record Commerce, BIS, or White House statement or published rule; the text of the Lutnick letter itself; original Wall Street Journal, Semafor, and Politico reporting retrieved directly rather than secondhand; on-the-record UK government or UK AISI statements given the institute's role in the disclosed caveat; Amazon's formal position beyond its non-confirmation; the full scope of allied-government and congressional reaction; rigorous legal analysis of the statutory basis and whether software-as-export precedent applies; and, above all, any resolution or material development after the date of this writing.

International and Political Reaction

The directive's reach to foreign nationals produced immediate effects outside the United States, and the early reaction abroad has framed the episode as much around sovereignty as around the specific safeguard dispute. British reporting noted that organizations in the United Kingdom that had been using the models—described as including hospitals and companies engaged in research—lost access overnight when the global shutdown took effect [16][18]. Conservative MP Tom Tugendhat framed the cutoff in geopolitical terms, arguing that "disabling Fable 5 and other models for foreigners is not a misunderstanding or a mistake, it's the inevitable result of technology shaping warfare so that sovereignty is more about code than cannons" [16]. *(This is a single legislator's characterization, attributed.)*

European commentary characterized the episode as a "wake-up call" on dependence for advanced AI capability on a single foreign supplier subject to that supplier's home-government controls [19]. These reactions are reported reaction and opinion, not established fact about the directive's terms; they are included here because allied-government and market response is itself part of what is currently known and is among the developments worth following.

The Open Letter

On June 14, 2026, a group of information-security leaders published the "Open Letter on Transparent AI Cyber Protections" at freefable.org. This section describes it factually; consistent with this document's editorial constraints, it offers the reader no advice about the letter.

The letter makes two principal asks. It calls on the U.S. government to lift the export-control directives on the Fable and Mythos models, and it sets out four principles it argues should govern future AI risk regulation: that such regulation be grounded in scientific evaluations developed with industry and academic input; that it be created through a democratic rule-making process; that it be enforced transparently and fairly, with time to remediate; and that it be applied only to the minimal extent necessary to ensure public safety [8].

Its empirical claims align with the arguments Anthropic has made on its own behalf. The signatories contend that while Mythos-class models are capable at finding software flaws and writing exploits, they are "not *uniquely* good" at these tasks, and that security professionals regularly use other foundation and open-source models—the letter names GPT-5.5, Opus, Sonnet, and Chinese models such as Kimi 2.7—for audits and red-teaming. They assert that Anthropic built multiple protections into Fable against offensive cyber use, characterizing those protections at launch as aggressive. And they argue that Chinese open-weight models remain only months behind the best American models, such that restricting one U.S. product does little to change the global capability picture [8].

Two observations about the letter, offered as factual characterization rather than endorsement, are relevant to a reader weighing it. First, the letter's empirical claims substantially track Anthropic's own position in the dispute. Second, its procedural asks resemble principles Anthropic itself has articulated—the company's statement calls for a statutory process that is "transparent, fair, clear, and grounded in technical facts" [1] [13]. The signatories include security-industry chief executives and founders, chief information security officers from major technology firms, and academic and former-government figures; the current signatory count is available on the letter's page and continues to change [8]. The alignment between the signatories' position and one party's interest is a fact a reader may wish to weigh; this document neither discounts the letter on that basis nor urges any response to it.

What to Watch

This is a fast-moving story, and the most useful posture for a reader is to track a small number of concrete developments that would materially change the picture.

The first is any official, on-the-record statement from the Commerce Department, BIS, or the White House, or any published rule in the Federal Register—any of which would, for the first time, allow the public to evaluate the government's reasoning against a citable document rather than inference. The second is the surfacing of the Lutnick letter's actual text, which would settle several open questions about scope and whether partial availability pending licenses was ever permissible. The third is any on-the-record statement from the UK government or UK AISI, whose red-team work appears connected to the evidence the government conveyed. The fourth is Amazon's formal position, if it offers one beyond its current non-confirmation. The fifth is the breadth of allied-government and congressional response, which will indicate

whether this is treated as an isolated product action or a precedent. The sixth is serious legal analysis—and potentially litigation—testing the statutory basis for restricting commercial AI model access on national-security grounds, an area that current commentary describes as not yet defined by statute, regulation, or court decision [13]. And the seventh, most consequential of all, is any resolution: a remediation accepted by the government, a license granted, or access restored, along with the terms on which any of that occurs.

The Cloud Security Alliance's interest in this episode is institutional rather than partisan. The questions it raises—how export-control doctrine designed for physical goods and technical data applies to cloud-delivered AI models, how deemed-export obligations scale on consumer platforms, and what a transparent, evidence-grounded process for adjudicating model-level cyber risk would look like—bear directly on the governance frameworks the AI security community is building. CSA will continue to follow the matter and will update this account as the record develops.

The single call to action this document offers is to keep learning. The facts established as of June 15, 2026, are documented above; the contested and unverified claims are flagged as such; and the developments most likely to clarify the record are listed here. Readers are invited to follow those developments and to revisit the epistemic labels as new evidence arrives.

References

Sources are rated by tier, reflecting proximity to primary evidence and editorial reliability rather than agreement with any side. Tier 1 sources are primary or first-party artifacts of the dispute; Tier 2 are established news organizations conducting original reporting; Tier 3 are trade press and secondary analysis that is more interpretive or single-sourced. A one-line note records what each source contributes and any conflict of interest.

Tier 1 – Primary / first-party

[1] Anthropic. "Statement on the US government directive to suspend access to Fable 5 and Mythos 5." June 12, 2026. – The company's own account; source for scope, the global-disable rationale, the "narrow/non-universal/benign" characterization, the GPT-5.5 comparison, and the "misunderstanding" framing. *First-party; a party to the dispute.*

[2] Anthropic. "Claude Fable 5 and Mythos 5." June 9, 2026. – Pre-dispute launch post; source for the 1,000+ hours testing claim, the explicit UK AISI caveat, and the under-5%-of-sessions routing to Opus 4.8. *First-party.*

[8] freefable.org. "Open Letter on Transparent AI Cyber Protections." June 14, 2026. – The security-executive petition; source for the asks, the empirical claims, and the (continually updating) signatory list. *Advocacy document; position aligns with Anthropic's.*

[9] Anthropic safeguards technical document (PDF). "Fable 5 / Mythos 5 safeguards." Referenced by the open letter. *First-party technical material.*

[6] David Sacks. "Thread on the Anthropic export controls." X, June 13, 2026. – Senior administration adviser's personal account of the government side, including the "refused to fix" claim. *First-party administration voice; personal social-media statement.*

[14] U.S. Government. "15 CFR 734.13 – Export (deemed export rule)." Code of Federal Regulations. – Primary regulatory text underpinning the deemed-export analysis in this paper. *Primary law; doctrinal background, not specific to this case.*

Tier 2 – Established news organizations (original reporting)

[3] Axios. "Trump admin blocks foreign access to Anthropic's most powerful AI." June 12, 2026. – Scoop on the Commerce letter's mechanics. *(Direct fetch was blocked at time of writing; cited via corroborating coverage.)*

[10] Axios. "Trump admin blocks foreign access to Anthropic's most powerful AI." June 12, 2026. — The scoop detailing the Commerce letter's licensing mechanics: a license required for export, re-export, or domestic transfer, with the trigger attributed to "another company" claiming it jailbroke Mythos. *Same article as [3]; cited here specifically for the mechanics.* See also Axios, "How Amazon and the White House ended Anthropic's Fable," June 13, 2026.

[5] NBC News. "Anthropic suspends new AI models Fable, Mythos after government directive." June 12, 2026. — Source for the Commerce Department's declining to comment.

[4] CNBC. "Anthropic disables access to Fable 5 and Mythos 5 to comply with government directive." June 12, 2026. — Corroborates timeline, Lutnick signature, and disproportionality argument.

[11] Fortune. "How a warning from Amazon led the White House to shut down Anthropic's Mythos model." June 14, 2026. — Reconstruction of Amazon's reported role and the "90 minutes" detail. *Relies substantially on anonymous sourcing; omits that Amazon is a major Anthropic investor.*

[12] Fortune. "Anthropic disables Fable and Mythos AI models following U.S. government export ban." June 13, 2026. — Corroborates scope and global-disable rationale.

[15] Fox Business. "Trump admin says Anthropic's 'recklessness' triggered export controls on latest AI models." June 2026. — Source for the anonymous-official "recklessness" framing. *Single anonymous official.*

[17] Semafor. "White House move to limit Anthropic linked to concerns about Chinese access to Mythos." June 13, 2026. — Source for the (hedged, anonymous) Chinese-access suspicion and the Amazon role. *Entirely anonymous sourcing; claims explicitly qualified in body text.*

[18] Time. "Anthropic Pulls Its Most Powerful AI Models After U.S. Bars Foreign Access." June 13, 2026. — Corroborates the foreign-access bar and global effect.

[16] The New Stack. "Federal government orders Anthropic to pull Fable 5 and Mythos 5, three days after launch." June 2026. — Source for the UK AISI red-team-lead context, the Tugendhat quote, and UK organizations losing access.

Tier 3 — Trade press / secondary analysis

[7] Benzinga. "Sacks: US Wants To Lift Anthropic Export Controls 'As Soon as Possible' After Fable 5 Fix." June 2026. — Source for Sacks's "as soon as possible" and cooperation quotes. *Secondary aggregation of Sacks's statements.*

[13] Volkov Law. "When the Government Pulls the Plug: Anthropic, Export Controls, and the Future of AI Governance." June 2026. — Legal/compliance analysis of the statutory-authority question and the undefined scope of BIS authority over AI model access. *One law firm's interpretation, not a judicial determination.*

[19] Euronews. "'Wake-up call': Europe reacts to Anthropic halting access to its Fable 5 and Mythos 5 AI models." June 13, 2026. — Source for European sovereignty reaction. *Reaction/opinion reporting.*

[20] Tom's Hardware. "Trump adviser David Sacks says Anthropic refused to fix Fable 5 jailbreak before US export controls." June 2026. — Trade-press summary of the Sacks "refused to fix" account and the (unverified) Chinese-access headline frame. *Interpretive; headline frame not substantiated in body.*

[21] Dark Reading. "Claude Fable 5 Doesn't Change the Mythos Security Story." June 2026. — Quotes the launch blog on the UK AISI caveat. *Secondary analysis.*

[22] Amazon. "Amazon and Anthropic deepen their strategic collaboration." aboutamazon.com. — Primary-source documentation of Amazon's multibillion-dollar investment position in Anthropic, establishing the partner-and-competitor relationship relevant to weighing the Amazon-warning reporting. *First-party Amazon announcement.*

Transparency note: This whitepaper was prepared with the assistance of an AI system built by Anthropic, which is a party to the dispute it describes. The source set includes Anthropic's own framing of events, which has been labeled as first-party and held to the same skeptical standard as every other party's account. Where the only available source for a detail is a party's statement about itself, that limitation is noted in the text.

Editorial note: This document makes no recommendation. Its sole call to action is an invitation to continue following the matter as the record develops. Epistemic labels reflect the state of the evidence as of June 15, 2026, and should be revised as new facts emerge.