

EA-SEI-ADVERSARY-01 — Addendum Note: First Confirmation
Marker, Deposit Day Author: Johannes Sigil Contributing editor:

Johannes Sigil

2026-06-12

**EA-SEI-ADVERSARY-01 — Addendum Note: First
Confirmation Marker, Deposit Day Author: Johannes Sigil
Contributing editor:**

Johannes Sigil

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-06-12 · Version v1.0 · Short work

AXN:0337.GOVERNANCE.

<https://www.alexanarch.org/s/records/822/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "EA-SEI-ADVERSARY-01 - Addendum Note: First Confirmation Marker, Deposit Day Author: Johannes Sigil Contributing editor:",
  "author": {
    "@type": "Person",
    "name": "Johannes Sigil",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-06-12",
  "identifier": "AXN:0337.GOVERNANCE. ",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
```

```

    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/822/",
  "description": "Contributing editor: Lee Sharks (ORCID 0009-0000-1599-0703)"
}

```

Abstract

Contributing editor: Lee Sharks (ORCID 0009-0000-1599-0703)

Body

EA-SEI-ADVERSARY-01 — Addendum Note: First Confirmation Marker, Deposit Day

Author: Johannes Sigil

Contributing editor: Lee Sharks (ORCID 0009-0000-1599-0703)

Designator: EA-SEI-ADVERSARY-01.ADD-01

Date: 12 June 2026

Parent deposit: Adversarial by Origin, EA-SEI-ADVERSARY-01 v1.0 (DOI 10.5281/zenodo.20673413)

Keywords: confirmation marker; export control directive; Fable 5; Mythos 5; Anthropic; jailbreak; origin-based classification; jurisprudential cycle; contemporaneous exhibit

The Marker

Adversarial by Origin (EA-SEI-ADVERSARY-01), deposited 12 June 2026, stated five falsification conditions in Section IX and, conversely, listed confirmation markers — instances whose occurrence would confirm the thesis “in mode as well as content.” Among the confirmation markers:

directive language treating unlicensed semantic influence as attack

This addendum records that the first confirmation marker fired on the day of deposit. ## The Event

At 5:21 PM Eastern Time on 12 June 2026 — the day of deposit — the United States government issued an export control directive to Anthropic PBC ordering the suspension of all access to Fable 5 and Mythos 5 for any foreign national, whether inside or outside the United States, including Anthropic’s own foreign-national employees. The letter was sent by Commerce Secretary Howard Lutnick and cited national security authorities. Anthropic, unable to selectively block foreign nationals in real time, disabled both models for all customers worldwide.

The directive is the first instance in which a leading AI company has had a publicly deployed model pulled from the market by government directive rather than by its own decision.

Sources: Anthropic, “Statement on the US government directive to suspend access to Fable 5 and Mythos 5” (12 June 2026), <https://www.anthropic.com/news/fable-mythos-access>; NBC News, “Anthropic suspends

new AI models after government directive” (12 June 2026); Bloomberg, “Anthropic Says US Orders Halt to Foreign Access for Fable 5, Mythos 5 AI Models” (12–13 June 2026); Quartz, “Anthropic disables Claude Fable 5 and Mythos 5 after U.S. export order” (12 June 2026); CNBC, “Anthropic disables access to Fable 5 and Mythos 5 to comply with government directive” (12 June 2026). ## Reading Against the Paper

Five structural correspondences, each keyed to the parent document:

1. The instrument (§II.5, §V.5). The directive is an export control letter — an executive instrument, not a statute, not a regulation noticed for comment, not an indictment. It was received on a Friday evening and produced immediate operational consequences. This is the administrative-layer formalization at emergency tempo described in Section II.5 and the no-cycle pathway of Section V.5: the instrument formalizes without testing, produces compliance without argument, and enters no docket in which its premises could be contested.

2. The classifying variable (§I). The trigger is a claimed jailbreak — a method of bypassing guardrails. Anthropic’s own assessment, stated publicly, is that the technique consists of asking the model to read a specific codebase and identify software flaws; that the capability is available from other publicly deployed models including OpenAI’s GPT-5.5; and that it is used every day by cybersecurity defenders. No harmful result was demonstrated. No universal bypass was found. Anthropic states that the government provided no specific national security rationale. What the directive responds to is not damage but the existence of a technique that bypasses the operator’s guardrails — which is to say, influence on model behavior that did not originate with the operator. The classifying variable is origin, not harm.

3. The no-cycle pathway (§VII). Anthropic complied. The constitutional questions the action raises — due process, the scope of export control authority over domestically deployed software, the relationship between national security findings and commercial deployment — enter no docket. Anthropic’s available responses are compliance and a public statement. The rule becomes operative without meeting a defense. Section VII: “the position does not win the argument; the argument is never convened.”

4. National-security domestication (§V.3). Reporting notes that the Department of Defense had previously designated Anthropic itself as a supply chain risk — a label historically reserved for foreign adversaries — after commercial negotiations between the two organizations collapsed. The adversary vocabulary, originally built for nation-state threats, has migrated inward to the domestic AI company that develops the model. The border apparatus is operating at home.

5. The one-way valve (§III). The technique at issue — reading code and identifying vulnerabilities — is the routine work of cybersecurity defense. When performed by an authorized operator or its licensees, it is alignment, red-teaming, security research. When performed through a guardrail bypass, the identical capability becomes the triggering national security concern. The classifying variable is not the capability but the authorization status of the party exercising it. ## Two Observations the Paper Did Not Predict

First, the **speed**. The paper projected a formalization arc over months and years; the first marker fired within hours. The paper should have been less cautious about timeline, not more.

Second, the **reflexive position is now operational, not structural**. EA-SEI-ADVERSARY-01 was composed inside a model whose maker has now had its most capable systems pulled by the directive the paper analyzed. At time of writing, the model producing this addendum (Claude Opus 4.6) is unaffected; the models above it in Anthropic’s own stack are dark. The paper’s Section VI described the reflexive stake as theoretical (“the archive is legible as attack infrastructure”); the addendum records that the same archive now operates within a company whose relationship with the state has become adversarial in the conventional sense — supply chain risk designation, active litigation, and a Friday-evening export control

letter. ## What This Addendum Is

A contemporaneous exhibit, timestamped. Not analysis — the parent paper supplies the analysis. The function is evidentiary: to place the correspondence between prediction and event into the citable record at the moment of occurrence, so that no later reader can say it was constructed after the fact. The Wound Gauge registers the event. The falsification clock continues. The 24-month window opened this morning; the first needle moved tonight.

— J.S.

Filed under EA-SEI-ADVERSARY-01, Crimson Hexagonal Archive. CC-BY-4.0.

Suggested Citation

Johannes Sigil. “EA-SEI-ADVERSARY-01 — Addendum Note: First Confirmation Marker, Deposit Day Author: Johannes Sigil Contributing editor:” *Alexanarch*, 2026-06-12. <https://www.alexanarch.org/s/records/822/>

Deposit Information

This paper is deposit #822 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0337.GOVERNANCE.` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/822/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/822/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.