NPSI / TECHNICAL SERIES / WORKING PAPER · 26 JULY 2026 · UNCLASSIFIED

WP-10

Fair Use for We, IP Theft for Thee

Nobody agrees what “open source” means in artificial intelligence. Washington is now writing policy inside that vacuum.

Jesse James · North Pacific Strategy Initiative
Working Papers on Pacific Sovereignty & Bilateral Architecture · Vol. I · Est. MMXXVI
jesse@npsi.ca · linkedin.com/in/jessecares

EXECUTIVE SUMMARY

Between 16 and 24 July 2026 the United States government came closer to restricting foreign open-weight AI models than at any point in the technology's existence. The Treasury Secretary put sanctions and Entity List designations on the table. The Director of the Office of Science and Technology Policy named a specific Chinese laboratory and a specific American model. Axios reported an administration actively weighing a ban.

Not one of those statements rests on an agreed definition of the thing being restricted.

“Open source” has no enforced meaning in artificial intelligence. The Open Source Initiative published a formal definition in October 2024. The largest company claiming the label rejected it. No regulator or procurement authority anywhere has adopted it as binding. In the vacuum, the word has become a marketing asset, a regulatory exemption, and now the object of a federal policy fight — while the licence files it supposedly describes say something else entirely.

This paper reads the licences. It finds that of eleven major model families surveyed, three satisfy the OSI definition, and none of those three operates at frontier scale. It finds that Meta's Llama 4 licence withholds rights from anyone domiciled in the European Union, inside a product marketed as open source, in the one jurisdiction that offers open-source models a regulatory exemption. And it finds that the characterisation now driving policy — that training on another party's output is fair use when American labs do it and theft when Chinese labs do it — is real at the level of framing, weaker than advertised at the level of evidence, and being applied to conduct that three separate bodies of law treat three separate ways.

Key findings

01 · DEFINITION

The OSI's Open Source AI Definition v1.0 has existed since October 2024 and has been adopted as a binding qualifying criterion by no regulator or public procurement authority located in this survey.

02 · LICENCE TEXT

Llama 4's licence states that rights are *not* granted to individuals domiciled in, or companies with a principal place of business in, the European Union.

03 · SCALE

Of eleven major model families surveyed, three meet the OSI definition — OLMo 2, Pythia/GPT-NeoX, and, arguably, BLOOM. All three sit well below frontier scale.

04 · LITIGATION

The largest copyright settlement in US history — $1.5bn, final approval 20 July 2026 — turned on how books were *acquired*, not on whether training is fair use. On the training question the same court ruled the other way.

05 · EXPOSURE

If Chinese open weights were declared legally tainted, Cursor's Composer 2 and Thinking Machines' Inkling are both downstream by their own published account.

What to do with this

If you ship on open weights: resolve the licence per model version, never per vendor. Mistral and Qwen both mix licences inside their own lineups. The tier table in §2 is built to be checked, not trusted.

If you write procurement rules: stop using “open source” as a qualifying criterion until your instrument defines it. Canada's Treasury Board directive expresses an open-source preference and supplies no definition at all.

If you write policy: the debate merges three legally distinct acts — taking weights, training on scraped outputs, and acquiring pirated corpora. Collapsing them is how both sides of this argument get to be right at once.

HOW TO READ THE EVIDENCE MARKS

Every load-bearing claim in this paper carries a grade. The grade is not a hedge; it is the reason to trust the document. A paper that displays a weak grade beside its most quotable allegation is making a harder promise than one that writes *reportedly* and moves on.

▉	**VERIFIED**	primary source retrieved and quoted directly
▓	**CORROBORATED**	two or more independent secondary sources agree
▒	**ATTRIBUTED**	on record from one interested party, uncorroborated
░	**UNDETERMINED**	could not be established; stated as a gap, never filled

Typography carries the same discipline. Serif is this paper's argument. Monospace is somebody else's words — licence text, quoted statements, docket numbers, figures. The reader never has to guess which is which, and the author cannot quietly blur the line.

CONTENTS

§1 The word nobody owns

§2 What the licences actually say

§3 Brussels wrote the incentive

§4 Three things the debate keeps merging

§5 What the courts have actually held

§6 The asymmetry, graded honestly

§7 Everyone is downstream

§8 The Canadian position

§9 What a fair reading concedes

§10 What could not be determined

§1 · THE WORD NOBODY OWNS

The Open Source Initiative spent two years building a definition, and the industry's largest claimant of the label declined to accept it.

OSAID v1.0 was published in October 2024 at All Things Open in Raleigh.[1] It sets out four freedoms: use the system for any purpose without permission; study how it works; modify it, including to change its output; and share it, with or without modification. Two of those freedoms have sharp edges. A field-of-use restriction breaks the first. A licence that revokes rights above a usage threshold breaks the fourth.

The definition made one large compromise. It requires *data information* — a description of the training corpus sufficient for a skilled person to recreate a comparable system — rather than the corpus itself, on the grounds that training data is encumbered by copyright, contract and privacy law in ways source code is not. The Free Software Foundation and the Software Freedom Conservancy objected publicly. The definition is contested from the purist side, not only the corporate one.[2]

Meta rejected it outright and continued to market Llama as open source. The practice has a name in the standards community: open washing.

> WHY THIS MATTERS NOW
>
> A word with no enforced meaning became load-bearing for three separate systems at once: a European regulatory exemption, a Canadian procurement preference, and an American national-security argument. None of the three supplies its own definition, and none of them points at OSAID.

§2 · WHAT THE LICENCES ACTUALLY SAY

Marketing language and licence text disagree, and the disagreement is checkable in about ten minutes per model.

Major model families by openness tier, with licence name and the specific restricting clause. Resolved per model version, not per vendor — Mistral and Qwen both mix licences inside their own lineups. Two rows are flagged: those are the ones where the marketing and the licence text say different things.

MODEL VERSION	TIER	LICENCE	RESTRICTING CLAUSE
OLMo 2 (AI2)	1 · OSAID	Apache-2.0	none material — weights, data, training code, checkpoints all published
Pythia / GPT-NeoX	1 · OSAID	Apache-2.0	none material — full pipeline published
BLOOM	1-adjacent	BigScience OpenRAIL-M	behavioural-use restrictions break freedom 1
DeepSeek V3 / R1	2 · weights	MIT	permissive weights; no training data, no training code
Kimi K2.5	2 · weights	Modified MIT	MAU display condition inside the modified terms
Kimi K3	2 · pending	unpublished	marketed as the largest open-weight model ever; weights withheld at the 16 Jul launch, scheduled 27 Jul. Licence terms not assertable until the card ships.
Qwen — varies	2 / 3	Apache-2.0 / bespoke	lineup mixes licences; resolve per version
Mistral — varies	2 / 3	Apache-2.0 / MRL	Research Licence bars commercial use on some versions
gpt-oss (OpenAI)	2 · weights	Apache-2.0	weights only
Llama 4 (Meta)	3 · restricted	Llama 4 Community	rights “not being granted to you if you are an individual domiciled in, or a company with a principal place of business in, the European Union”; 700M-MAU revocation; naming requirement
Gemma	3 · restricted	Gemma Terms of Use	prohibited-use schedule
Claude · GPT · Gemini	4 · closed	proprietary	weights not distributed

▮ VERIFIED — licence texts retrieved directly. ▒ UNDETERMINED — Kimi K3's licence, which had not published at the time of writing.

Three families clear the OSI bar. None of them competes at the frontier. Every model a developer would actually reach for in production sits in tier 2 or below, which means the practical question is never “is this open source” but “which specific clause applies to me.”

§3 · BRUSSELS WROTE THE INCENTIVE

The EU AI Act offers open-source models a partial exemption, defines the term itself, and does not use OSI's definition.

Article 53(2) relieves providers of models released under a free and open-source licence from the technical-documentation duty and the downstream-information duty. It does *not* relieve them of the copyright-policy duty or the training-content-summary duty. The exemption is partial, and it evaporates entirely for models with systemic risk — presumed above 10^{25} FLOPs — at which point the full obligations of Articles 53 and 55 attach regardless of licence.[3]

The Act supplies its own definition: a licence that allows access, usage, modification and distribution of the model, with parameters including weights, architecture information and usage information made publicly available.

> THE IRONY, STATED PLAINLY
>
> Llama's licence excludes EU-domiciled persons from the grant. A licence that withholds distribution rights from an entire jurisdiction is unlikely to satisfy a definition built on "access, usage, modification, and distribution" — which means the product most aggressively marketed as open source is the one least likely to qualify for the exemption that made the label valuable.

This is the second act of the story, and it is the more durable one. The regulatory exemption created a reason to claim the label. The absence of an enforced definition made the claim costless. Everything downstream — including the policy fight now underway in Washington — inherits that vacuum.

§4 · THREE THINGS THE DEBATE KEEPS MERGING

Three legally distinct acts are being discussed as one. Separating them is the single most useful thing a reader can take from this paper.

> ONE · TAKING MODEL WEIGHTS
>
> Theft. Uncontested. Not what is alleged against any Chinese laboratory in the present dispute.

> TWO · TRAINING ON SCRAPED OUTPUTS
>
> Contested and actively litigated. The copyright status of model outputs is unsettled. This is what Anthropic alleges against DeepSeek, Moonshot and MiniMax — and it is also what Cursor and Thinking Machines did to Kimi, and what Musk has testified xAI did to OpenAI.

> THREE · ACQUIRING PIRATED CORPORA
>
> Already adjudicated, and adjudicated against the lab. This is the actual basis of the $1.5bn Anthropic settlement. It is a question about *how you obtained the books*, not about what training does.

The rhetorical move available to every party is to describe conduct in category two, add a terms-of-service violation, borrow the word “theft” from category one, and validate the whole package by pointing at a settlement from category three. Each step is individually defensible. The assembled claim is not.

§5 · WHAT THE COURTS HAVE ACTUALLY HELD

American courts have drawn a consistent line: acquisition can be infringement, training generally has not been, and fair use collapses when the output substitutes for the source in its own market.

Bartz v. Anthropic

VERIFIED Final approval came on 20 July 2026 from Judge Araceli Martínez-Olguín in the Northern District of California. The class comprises 482,460 eligible works with 440,490 claimed, a 91.3% rate, at roughly $3,000 per work. Counsel fees were cut from the $187.5m requested to approximately $101.6m; all 54 objections were overruled; service awards were trimmed from $50,000 to $15,000 per class representative.[4]

The predicate matters more than the number. In June 2025 Judge William Alsup held that training on lawfully acquired books is fair use, and that downloading roughly seven million pirated copies from LibGen and PiLiMi was not. The settlement resolves the second holding. The first still stands, and Anthropic has said so:

> "We reached this settlement in 2025, after the court's landmark ruling that training AI on books is fair use under copyright law, which remains the law today."
>
> Aparna Sridhar, Deputy General Counsel, Anthropic · reported 20 Jul 2026

Thomson Reuters v. Ross Intelligence

VERIFIED Judge Stephanos Bibas granted partial summary judgment to Thomson Reuters on 11 February 2025, rejecting fair use as a matter of law over 2,243 Westlaw headnotes. Two questions were certified to the Third Circuit on 23 May 2025 under 28 U.S.C. §1292(b). The reasoning is the strongest doctrinal thread in this entire dispute and the least discussed: the fourth factor, market effect, was treated as most important, and the court found Ross meant to compete with Westlaw by building a market substitute — harming both the legal-research market and a potential derivative market for AI training data.[5]

> "Copying protectable expression to create a competing substitute isn't innovation: it's theft."
>
> Thomson Reuters, Third Circuit brief · November 2025

That sentence is the entire American AI industry's exposure written by the other side. It is also, read carefully, the argument against distillation — and against training on publishers, in exactly equal measure. Which is the problem.

The rest of the map

VERIFIED *NYT v. OpenAI* remains in discovery before Judge Sidney Stein in the Southern District of New York; the Times seeks model destruction under 17 U.S.C. §503(b). In the United Kingdom, *Getty Images v. Stability AI* [2025] EWHC 2863 (Ch) ended on 4 November 2025 with Getty abandoning its primary copyright and database claims mid-trial and losing on secondary infringement, retaining only what the judgment itself described as extremely limited trademark findings.

§6 · THE ASYMMETRY, GRADED HONESTLY

The characterisation asymmetry is real. The conspiracy version of it is not established, and this paper will not assert it.

Anthropic's public post of 23 February 2026 is the primary text. It was retrieved in full. It says:

> "We have identified industrial-scale campaigns by three AI laboratories—DeepSeek, Moonshot, and MiniMax—to illicitly extract Claude's capabilities to improve their own models. These labs generated over 16 million exchanges with Claude through approximately 24,000 fraudulent accounts, in violation of our terms of service and regional access restrictions."
>
> anthropic.com · 23 Feb 2026 · retrieved in full · figures are self-reported and not independently verified

▌VERIFIED The phrase "IP theft" does not appear. The document's vocabulary is *illicit*, *fraudulent*, *in violation of our terms of service*. The harm framing is national security — stripped safeguards, export-control circumvention — not copyright. Whether by design or by instinct, this is the language of a party that has a live fair-use position to protect.

Officials were less careful. ▌VERIFIED

> "We support open-source AI and the innovation it unlocks. But open source is not open season on American IP. When PRC firms conduct covert, industrial-scale distillation attacks that cross the line into IP theft, sanctions and Entity List designations will be on the table."
>
> Scott Bessent, Secretary of the Treasury · X · July 2026

> "We have information that Moonshot AI distilled Anthropic's Fable for the development of its K3 model. To do this they developed a sophisticated internal platform to conduct large scale distillation against U.S. models, allowing them to quickly switch between multiple methods of access to avoid detection."
>
> Michael Kratsios, Director, OSTP · X · 22 Jul 2026

So the shape of the asymmetry is documented. A private party avoids the copyright term while government officials adopt it, and an Anthropic policy staffer publicly endorsed the Kratsios framing.

> ▒ ATTRIBUTED — AND THIS IS WHERE THE PAPER STOPS
>
> David Sacks has alleged on the All-In Podcast that Anthropic made the IP-theft argument to government directly while withholding it in public. That allegation traces to one interested party — a sitting White House AI czar who is also a venture investor with portfolio exposure to the outcome — and no named outlet corroborates it. It is recorded here as an allegation. It is not a finding, and nothing in this paper depends on it.

The subtler reading is the more defensible one, and it is worse for everybody. Anthropic's wording was careful precisely because the trap is real. The actors in this dispute understand the asymmetry perfectly well. That is why the public language is drafted the way it is.

AUTHOR'S DISCLOSURE

On the conflict in this paper

This analysis was researched and drafted with substantial assistance from Claude, a model made by Anthropic — the party whose conduct §6 examines. The parts of the argument most favourable to Anthropic were adversarially tested, its strongest defence is stated in its own words above, and the single claim that would have been most damaging to it is the one graded ▒ ATTRIBUTED and excluded from the findings. Readers should weight the framing accordingly and check §5 and §6 against the primary sources, all of which are listed.

§7 · EVERYONE IS DOWNSTREAM

If Chinese open weights are declared legally tainted, a meaningful share of the American startup stack goes with them.

VERIFIED Thinking Machines' Inkling, launched around 15 July 2026, was by the company's own account bootstrapped with supervised fine-tuning on synthetic data generated by open-weight models including Kimi K2.5, with architecture referencing DeepSeek-V3.

VERIFIED Cursor's Composer 2 used Kimi K2.5 as its base and post-trained on top. The base was not disclosed at launch. Both facts come from Cursor:

> "It was a miss to not mention the Kimi base in our blog from the start."
>
> Aman Sanger, co-founder, Cursor · X

CORROBORATED Nearly 200 startups, organised through the newly formed Little Tech Association and including Y Combinator and Proton among its members, wrote to the President, Commerce, State, Treasury and OSTP urging policymakers not to conflate legitimate model-development techniques with misappropriation. A separate letter from Nvidia, Microsoft, Meta, Palantir and roughly twenty others warned against premature restrictions. Neither OpenAI nor Anthropic signed either letter.[6]

> THE UNANSWERED LEGAL QUESTION
>
> No court and no published legal analysis located in this survey addresses whether a model distilled from an allegedly infringing model is itself infringing. The derivative-taint question — the one that determines the actual blast radius of any ban — has never been tested.

There is also a technical objection to the government's causal story, and it comes from people with no obvious stake:

> "I don't think you get a model this strong and this quickly on the heels of Fable doing strictly distillation … Fable's only been publicly available since July 1st. You can't distill that much data, train a model, and release it in two weeks."
>
> Braden Hancock, Snorkel AI / Laude Institute · reported 23 Jul 2026

§8 · THE CANADIAN POSITION

Canada expresses an open-source preference in binding policy and never defines the term.

VERIFIED The Treasury Board Directive on Automated Decision-Making, in force since 2019 and substantially amended in April 2023, provides that departments should ensure, where possible, that the software components of automated decision systems are obtained under an open source licence, with source code published unless classified or statutorily exempt. It is a transparency preference, not a hard qualifying criterion, and it is not anchored to OSI, OSAID, or any other external definition.

The practical consequence for a Canadian public body is that the phrase does no work. A vendor can satisfy the preference with a licence that revokes above a usage threshold, restricts fields of use, or excludes a jurisdiction, because nothing in the instrument says otherwise. Layer on Quebec's Law 25, PIPEDA, ITSG-33 and the residency questions raised by the US CLOUD Act, and the operative decision — hosted API versus self-hosted weights — turns out to be governed by data-residency law rather than by openness at all.

> UNDETERMINED
>
> No Canadian federal, provincial or municipal body was identified as having deployed self-hosted open-weight models with country of origin as a documented deciding factor. The absence of evidence here is a gap in the public record, not a finding of absence.

§9 · WHAT A FAIR READING CONCEDES

Every counter-position in this dispute has a serious version. Here they are, unweakened.

Distillation is benchmarking, not theft. Weights are the software; outputs are not. Google submitted large volumes of queries to Yahoo and Microsoft search to benchmark its own rankings, and nobody concluded it had stolen an algorithm. Car manufacturers tear down competitors' vehicles as routine practice. Learning from observed output is an established and lawful form of competitive analysis.

The remedy is contractual, not statutory. If the conduct is a terms-of-service violation, enforce the terms of service. Implement know-your-customer at the API and at the data centre. Copyright law and trade law are extraordinarily blunt instruments for a problem that contract law already addresses.

Security review of foreign-origin weights is legitimate. The joint UK AISI and US CAISI evaluation of Kimi K3 is a real mechanism producing a real finding — in that instance, that the model's cyber capabilities sat significantly below other frontier models. Evaluating foreign models is not protectionism, and treating every security question as pretext is its own failure of analysis.

Frontier labs face genuine asymmetric exposure. Billions of dollars of training investment can be partially transferred through an API at the cost of inference. A company that spends the money has a real grievance against one that does not, whatever the correct legal characterisation turns out to be.

OSAID is contested from the purist side too. The FSF and Software Freedom Conservancy objections are not corporate lobbying. The definition's data compromise is a genuine weakness, and OSI itself describes v1.0 as the start of a conversation.

And the gap is smaller than this paper's framing implies. For nearly every practical deployer, the distance between an MIT-licensed weights file and full OSAID compliance is close to zero. You can run it, study it, modify it and share it. The missing training-data documentation matters for reproducibility, auditability and legal certainty — it does not usually matter on a Tuesday.

§10 · WHAT COULD NOT BE DETERMINED

Gaps are stated, never filled.

> ▒ UNDETERMINED
>
> Whether Anthropic made an IP-theft argument to government privately while avoiding the phrase publicly. Sourced to one interested party; no named outlet corroborates.

> ▒ UNDETERMINED
>
> Whether a model distilled from an allegedly infringing model is itself infringing. No court, no published legal analysis.

> ▒ UNDETERMINED
>
> Kimi K3's licence terms, which had not published at the time of writing. Weights were scheduled for 27 July 2026 and had not shipped as at 25 July.

> ▒ UNDETERMINED
>
> The current state of the Kimi-Researcher repository, whose June 2025 page promised to open-source both base and RL-trained models “in the following months.”

> ▒ UNDETERMINED
>
> Whether any Canadian public body has deployed self-hosted open-weight models with country of origin as a documented factor.

Two further notes on method. Every revenue and market-share figure circulating in the podcast coverage of this dispute was excluded, because the underlying estimates are third-party, unattributed and internally inconsistent — the economic-competitiveness argument is therefore absent from this paper rather than weakly sourced within it. And Kimi K3's arena results are blind human preference on arbitrary prompts across a preliminary vote count, not task performance; they are reported here as what they are and load-bear nothing.

SOURCES

Primary sources were retrieved directly wherever the mark ▮ appears. Where a claim rests on secondary reporting, the outlet is named in the text.

1 **Open Source Initiative.** Open Source AI Definition v1.0. Published October 2024, All Things Open, Raleigh. opensource.org

2 **Free Software Foundation; Software Freedom Conservancy.** Public objections to the OSAID data-information compromise, 2024–2025.

3 **Regulation (EU) 2024/1689 (AI Act).** Chapter V, Arts. 51–56; Art. 53(2) open-source exemption; Art. 54(6); Recitals 102–104.

4 **Bartz v. Anthropic PBC.** N.D. Cal. Final approval 20 July 2026, Martínez-Olguín J. Predicate ruling June 2025, Alsup J. Class figures per The Authors Guild.

5 **Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc.** D. Del., Bibas J., 11 February 2025; §1292(b) certification 23 May 2025; Third Circuit briefing November 2025.

6 **Little Tech Association** letter to the President, Commerce, State, Treasury and OSTP, July 2026, reported by Politico. Separate industry letter reported by CNBC and TechCrunch, 24 July 2026.

7 **Anthropic.** "Detecting and preventing distillation attacks," 23 February 2026. Letter to Sens. Scott and Warren, 10 June 2026, reported by CNBC 24 June 2026.

8 **Axios.** Maria Curi, "The secret Trump administration battle to fight Chinese AI," 20 July 2026; "White House draws new AI line on China," 24 July 2026.

9 **TechCrunch.** Thinking Machines Inkling launch, 15 July 2026; "Experts say exploiting Anthropic's Fable isn't how Kimi K3 got so good," 23 July 2026.

10 **Treasury Board of Canada Secretariat.** Directive on Automated Decision-Making, 2019, amended April 2023. Guide on the Use of Generative AI, June 2024.

11 **Getty Images (US) Inc & ors v Stability AI Ltd** [2025] EWHC 2863 (Ch), Joanna Smith DBE J, 4 November 2025.

12 **Meta.** Llama 4 Community Licence Agreement and Acceptable Use Policy. **Moonshot AI.** Kimi K2.5 / K2.6 model cards, HuggingFace.

NPSI WP-10 · Fair Use for We, IP Theft for Thee · 26 July 2026 · Unclassified · Open distribution.
North Pacific Strategy Initiative · Working Papers on Pacific Sovereignty & Bilateral Architecture · Vol. I · Est. MMXXVI
Jesse James · jesse@npsi.ca · linkedin.com/in/jessecares · npsi.ca

Corrections are welcomed and will be published. Every claim in this paper carries a grade; if you can move one, write.