SOFTWARE / MODELS / SILICON HARDWARE / HUMAN PROGRESS

Impact of Export Controls on the Future of Open Source AI

Vikrant Rathore

Why export controls reach beyond code and model weights into the physical hardware beneath open-source AI, and what happens to humanity when shared minds meet border walls.



23

AI & CHIP GATES

05

BIOMEDICINE

12

SOLAR MACHINES

ILLUSTRATED EDITION • 02 NOV 2026

MOFCOM 48 / US EAR

Impact of Export Controls on the Future of Open Source AI

Vikrant Rathore / CC BY-SA 4.0 Third-party materials retain their own copyrights and licenses. See NOTICE.

Jennifer Yuan, Vicky Ni, Ronak Rathore, Kanak Rathore
Version 2026-11-02
© 2026, Vikrant Rathore; original work CC BY-SA 4.0
2026-11-02
Built with the help of GuideDog, with assistance from Jennifer Yuan, Vicky Ni, Ronak Rathore, and Kanak Rathore.

Contents

01. The Day Knowledge Got a Customs Booth 9
- One Friday Afternoon, Two Very Different Clocks 9
- Crossing a Border Without Packing a Single Crate 10
- Four Words That Change Everything 11
- What This Means for the Two Teams 11

02. The Map of Forty Guarded Technologies 13
- How Governments Measure an Idea 13
- Reading the Tripwires Without Getting Tripped 14
- Artificial Intelligence, Chips & Superclusters (23 Checkpoints) 14
- Biomedicine & Gene Platforms (5 Checkpoints) 16
- Solar Silicon, Wafers & Factory Machinery (12 Checkpoints) 17
- What This Means for the Two Teams 17

03. Mind and Silicon: Guarding Models, Chips, and Clusters 19
- The Sourdough Starter, the Bakery Window, and the Cast-Iron Oven 19
- Three Tripwires for a Frontier Model (`266504X`) 20
- Data Red Lines (`266501J` & `266504X(1)`) and Autonomous Agents (`266505X` & `266502J`) 20
- The Hardware Layer Inside China's AI List: Chips, Memory, and 10,000-Card Clusters 21
- What This Means for the Two Teams 21

04. Open Models: Public Releases and Private Transfers 23
- An Open License Is One Part of a Release 23
- Start with the Actual Artifact 23
- China: Registration Now, Catalogue Proposal Separately 24
- United States: A Separate Jurisdiction Test 24
- Possible Effects on Open Models 24
- Independent Review and Its Limits 25
- A Practical Release Record 25
- What This Means for the Two Teams 25

05. Living Code: Gene Platforms and Shared Cures 27
- Guarding the Printing Press While Sharing the Books 27
- Undisclosed Biological Targets and Biosecurity Walls 28
- What This Means for the Two Teams 28

06. Capturing Sunlight Inside a Machine 29
- You Can Ship the Solar Panel, Not the Secret Kitchen 29
- Footnote 4: When the Physical Machine *Is* the Recipe 30
- What This Means for the Two Teams 30

07. Washington's Gatekeepers: Silicon Chips and Model Weights 31
- Guarding the Physical Engines of Global Computing 31
- The January 2026 Semiconductor Rule: Permission to Knock, Not an Open Door 31
- The Strange Legal Life of America's AI Model-Weight Rule (`4E091`) 32
- Two Neighboring Washington Gates: Venture Dollars and Bulk Sensitive Data 32
- What This Means for the Two Teams 32

08. Two Border Systems and How to Get Official Approval 33
- Two Philosophies Meeting Inside the Same Project 33
- Step-by-Step: How to Register or Apply for Official Approval 34

China: Government Departments, Statutory Timelines, and Verified Contacts . . . 35
United States: Federal Agencies, SNAP-R Workflow, and Verified Contacts . . . 36
What This Means for the Two Teams . . . 37

09. Everyday Situations: A Reader's Field Guide 39
Can I Still Download DeepSeek, Qwen, or Other Open Models onto My Own Hardware? . . . 39
Does Hardware Affect Open-Source AI Even When the Software License Is MIT or Apache 2.0? . . . 39
If Overseas Users Chat with Our AI App or Call Our API, Is That a Model Export? . . . 39
Our Lab Stays in China, but Foreign Investors Are Acquiring Control. Do We Register? . . . 39
What Should We Record in a Two-Page "Release Passport" Before Publishing? . . . 40
Six Real-World Scenarios . . . 40

10. Primary Sources, Method, and Plain Glossary 43
How Every Fact in This Book Was Verified . . . 43
Plain-Language Glossary for Normal Humans . . . 43
Complete Register of 26 Verified Primary Sources . . . 44
Research and Historical References Added in This Review . . . 45
Credits and Reuse . . . 46
Optional: Checking the Game Example . . . 46
Optional: Independent Review Is Not a Safety Guarantee . . . 46

An idea can be shared without using it up. Thomas Jefferson expressed this in an 1813 letter: *"He who lights his taper at mine, receives light without darkening me."* That image captures the promise of open-source AI: one team can learn from another without taking its knowledge away. JEFFERSON



Fig. 1 A border barrier holds up computer parts and technical plans while two research teams try to work together.

Open-weight models can give researchers and smaller organizations more control over how they use AI. A clinic, a language project or a university might adapt a public model on its own computers. These are illustrative uses, not documented deployments or proof that a model is clinically safe. Downloadable weights also differ from the fuller requirements of the Open Source AI Definition. OSI R1 QWEN

Sharing a model is only the beginning. To turn shared knowledge into a useful tool, a team also needs computers and a way to keep learning from other teams.

Not Just Software or Weights: The Hardware Beneath Open Minds

A model can travel as code and learned numbers. The computers that train and run it are physical machines. An open software license and access to those machines solve different parts of the problem.

Open models depend on physical computers. Large training runs can need advanced processors, memory and fast networks; smaller models can run on less demanding equipment. Not every open-model user needs the latest manufacturing process or a cluster of ten thousand cards. Controls on scarce hardware can nevertheless affect which teams can train, inspect and deploy models.

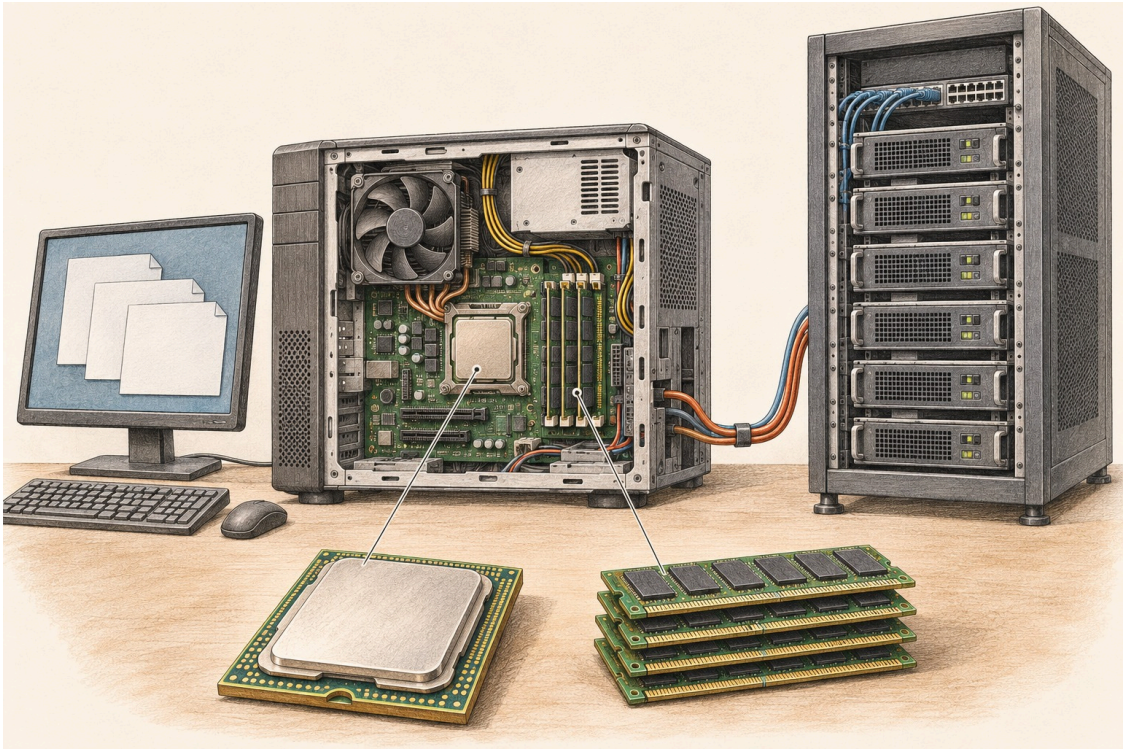


Fig. 2 Open models depend on chips, memory and computing infrastructure. A restriction on one part can affect access to the whole system.

- **Hardware and design knowledge:** US controls cover specified chips, manufacturing items and some foreign-made items under defined conditions. China proposes controls on specified processor and memory designs, and on certain technologies for clusters with **more than 10,000** accelerator cards. Design knowledge, finished chips and clusters are different objects of control. US26 EAR734 EAR742 CN-A1
- **Models and data:** Announcement 48 began temporary registration for covered transfers on **16 October 2026**. Its proposed catalogue includes model core technology meeting any of three tests: training calculation, parameter count or a qualifying global top-10 assessment. These proposed categories are separate from the registration already operating and from the US model-weight rule. CN48 CN-A1 US25 GAO26

> **IMPORTANT**
> A team may have permission to use an open model yet struggle to obtain suitable computers. Our story therefore follows both the knowledge being shared and the hardware needed to use it.

The Fortress Trap: How Both Sides Can Lose

The hardware problem leads to a political problem. A useful open model depends on contributions from many teams, yet each government also wants to protect its country's capabilities. How can choices made for national advantage end up holding back both sides? We can follow the problem through a fictional shared workbench.

Two Teams, One Shared Starting Point

Imagine a research team in Country A and another in Country B. Both start from the same publicly available AI model. They want to adapt it for ordinary local-language work. A's team finds a way

to run it with less memory. B's team develops tests that uncover errors and helps fix them. These are imagined teams and contributions, not accounts of actual deployments.

If the teams exchange their next improvements, each can build on the other's work. A receives better testing; B receives a model that needs fewer resources. Both still need computers, staff and time. Sharing knowledge helps them use those resources; it does not make physical hardware free.



Fig. 3 Two teams finish a computer by sharing parts and knowledge, then struggle when a wall separates them. This illustrates the assumed game, not a measured outcome.

Two Choices for Each Government

Now the governments make policy for the next round of work. Each wants useful progress at home. Each also worries about giving the other country an advantage without receiving anything in return. We simplify their choices to two:

- **Share:** Allow the teams' ordinary research improvements to be exchanged.
- **Wall:** Broadly restrict the country's new contributions from reaching the other side.

The wall concerns future contributions. It cannot make copies of an already-public model disappear. The teams can keep their starting files, but their next improvements may stop reaching each other. Real export controls offer many more choices, including narrow restrictions and specific permissions. We first set those complications aside so we can see one incentive problem clearly.

Four Possible Endings, Then Four Scores

We now give the outcomes invented preference scores. A higher score means that a government prefers that outcome in this story. The scores combine its assumed concern for useful progress, protection and relative advantage. They are not AI test results, percentages or measured national gains. Nothing in the announcement supplies these numbers.

Both share: 17 each. A's memory-saving improvement reaches B, and B's testing and fixes reach A. Both teams improve their work without independently repeating every discovery. We give each government 17 to represent a good result from mutual cooperation.

A shares while B builds a wall: A gets 3, B gets 20. B can use A's newly shared improvement while withholding its own new contribution. We assume B values that one-sided advantage most,

so it gets 20. A gets 3 because we assume its government dislikes sharing alone while the other country protects its contribution. A's team can still do useful work; 3 expresses its government's least-preferred outcome here.

A builds a wall while B shares: A gets 20, B gets 3. The roles reverse. A takes the one-sided advantage, while B is the only country sharing. We give both countries the same assumed priorities so we can see the strategic problem without assigning today's United States or China to either role.

Both build walls: 5 each. Neither receives the other's next contribution. They keep working from the public starting point, but must repeat more work or go without useful improvements. We assume each government prefers protecting its contribution to being the only side sharing, so 5 is above 3. We also assume both value mutual cooperation more than mutual isolation, so 17 is above 5.

The hardest assumption is that **20 is better than 17**: each government values a one-sided advantage more than shared progress in this round. That preference creates the temptation to restrict. If a government instead values mutual progress most, this particular trap may not arise. The numbers describe an assumed ranking; choosing them does not prove that real governments have it. GAME

Reading the Game Picture

Country A's choice selects a row. Country B's choice selects a column. The box where they meet gives A's score beside the letter A and B's score beside B. For example, the top-right box means A shares and B builds a wall: A gets 3 and B gets 20. The green box is mutual sharing; the red box is mutual restriction.

INVENTED PREFERENCES / HIGHER MEANS PREFERRED

How both sides can lose

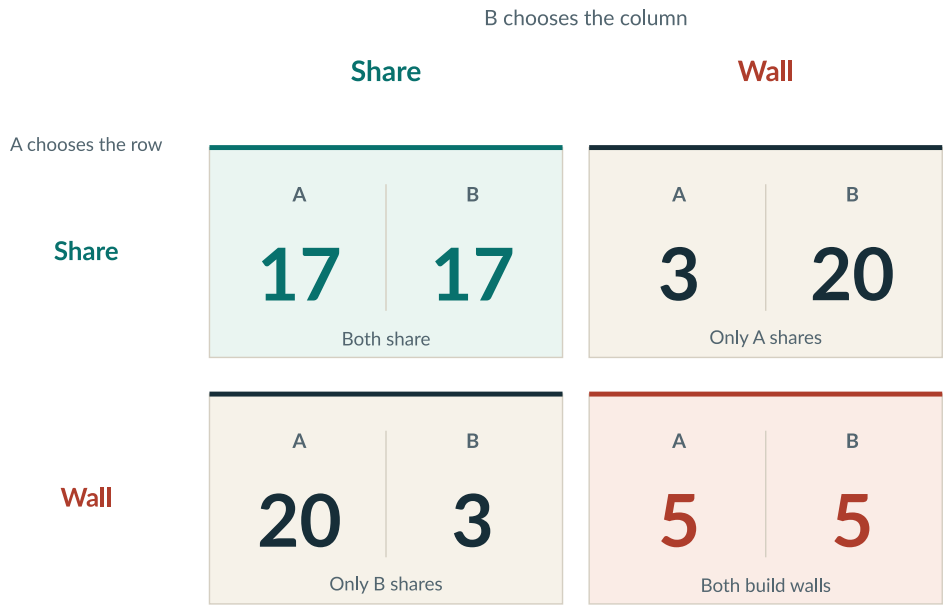


Fig. 4 Four endings for the two teams. The scores are invented preferences; higher means preferred within this example.

Why the Wall Starts to Look Like the Safer Choice

A's government considers what B might do. **If B shares**, A compares sharing for 17 with withholding its own contribution for 20. Under our assumptions, A chooses the wall to gain the one-sided advantage. In the picture, this means comparing the two boxes in the left column.

If B builds a wall, A compares sharing alone for 3 with also restricting for 5. A again chooses the wall, this time to avoid being the only side sharing. In the picture, this means comparing the two boxes in the right column.

B reasons the same way. Each government can tell itself that its wall is sensible whatever the other does. Both therefore choose walls and receive 5 each, even though both would prefer the 17 each available through cooperation. Neither has to want a poorer joint result for that result to emerge.

This is the **Prisoner's Dilemma**. The ending is also a **Nash equilibrium**: with the other wall still in place, neither government improves its score by changing its choice alone. For A, opening alone would move its score from 5 down to 3. Mutual sharing needs both to change, or a change in what makes cooperation worthwhile. GAME

What This Explains About Export Controls

The story puts a policy question behind the technical entries in the next chapters. A government may see protection or bargaining advantage in a broad restriction. Across both countries, the same restrictions may interrupt complementary improvements, duplicate effort and make smaller research teams' work harder. Those are mechanisms to investigate, not costs that our invented scores have measured.

Real US and Chinese controls are not this symmetric two-choice game. The countries have different capabilities and legal tools. A reporting duty, a prohibition, a normal license, a public-information exclusion and an easier release route can affect cooperation differently. The following chapters separate those cases. We do not label every filing requirement a wall or every exclusion a complete freedom to transfer.

A narrow control that prevents serious harm could also change the ranking of outcomes. Public model access has benefits and misuse risks. The policy task is to identify the specific risk being controlled and the cooperation being lost, rather than treating either complete openness or broad closure as an automatic answer. NTIA OPENRISK

Can the Teams Keep Working Together?

Return to the workbench. The teams expect to make another improvement next month, and another after that. A government that takes a short-term advantage today may lose access to future cooperation it values. A credible continuing arrangement can therefore change the choice. Whether that works depends on observable actions, the value of future exchanges and the possibility of mistakes; a promise alone does not solve the problem.

We carry one question into the rest of the book: **which rules address a specific danger while leaving a workable route for the next useful contribution?** As we examine model files, hardware, medical platforms and solar equipment, the object being shared changes. The question about lost cooperation remains. An optional calculation is in Optional: Checking the Game Example; the story can be followed without it.

The Central Factual Discovery: An Ajar Doorway, Not a Sealed Wall

The fictional wall hides distinctions that matter in the real rules. Is the technology already public? Is a new transfer prohibited, restricted or eligible for an easier route? Announcement 48 and the US rules answer different questions. We now turn from our imagined workbench to the source texts:

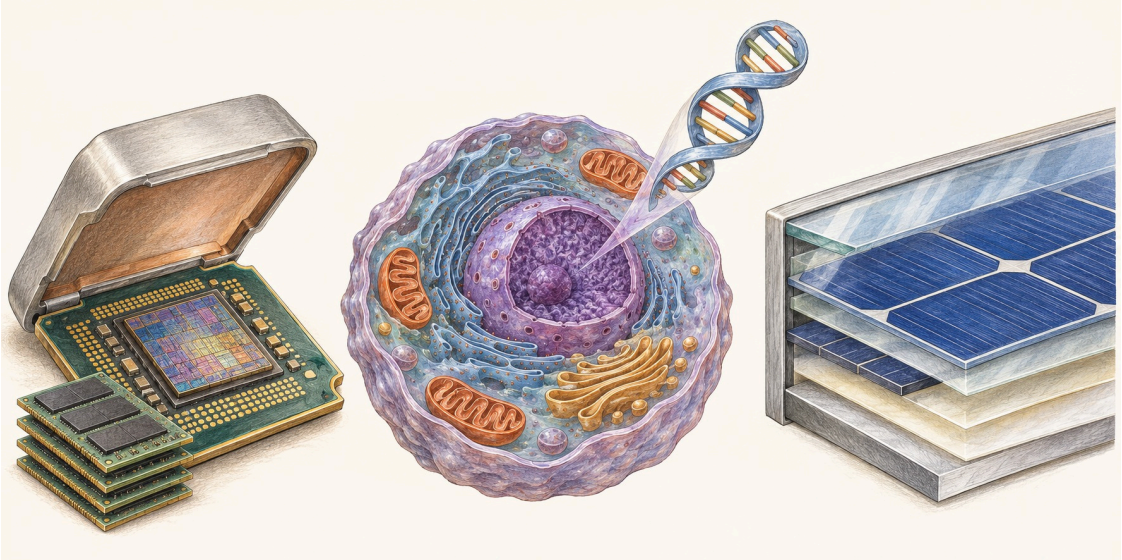


Fig. 5 Computing, biomedicine and solar manufacturing, shown through a processor, a cell and the layers of a solar panel.

Table 1 The 40 proposed technology points

Area	Proposed prohibited	Proposed restricted
AI and computing hardware	6	17
Biomedicine	3	2
Solar manufacturing	5	7

- **Already public and newly released are different.** The Chinese draft excludes AI technology already public through open source or papers. US rules have their own published-information provisions and model-weight notes. A public paper does not necessarily disclose private weights; a new private derivative needs its own assessment. Neither exclusion settles every other legal question. CN-A1, note 1 EAR734 EAR774
- **New open-source AI and chip tape-outs get a dedicated fast lane.** Rather than banning open releases of restricted AI, China's Footnote 2 proposes **automatic licensing and facilitation** ("apply and go" in MOFCOM's official explanation) for open-source exports, overseas foundry chip tape-outs (芯片流片), and Chinese firms' self-operated overseas compute clouds CN-A1, note 2 CNQA.
- **Registration is current; the new prohibitions remain proposed.** Covered ownership transfers, overseas licenses and foreign-control changes require attention now. The new core-data prohibition is in the draft, not an already-final ban created by Announcement 48. CN48 CN-A1

How to Read This Book

We wrote this book for normal, curious human beings (builders, researchers, doctors, founders, students, journalists, and citizens) who want vivid clarity backed by verified primary evidence:

- 01. The Day Knowledge Got a Customs Booth and 02. The Map of Forty Guarded Technologies unpack Beijing's Announcement 48, its two separate clocks, and all **40 proposed technology points**.
- 03. Mind and Silicon: Guarding Models, Chips, and Clusters takes you inside both the software recipe and the silicon hardware stack (`266504X`, `266506X`, `266507X`).
- 04. Open Models: Public Releases and Private Transfers distinguishes public releases from private transfers and keeps current duties separate from proposed routes.

- 05. Living Code: Gene Platforms and Shared Cures and 06. Capturing Sunlight Inside a Machine explore living gene platforms and solar machines where know-how travels inside physical equipment.
- 07. Washington's Gatekeepers: Silicon Chips and Model Weights demystifies Washington's chip (`3A090`) and model-weight (`4E091`) rules, including the May 2026 GAO clash USNE GAO26.
- 08. Two Border Systems and How to Get Official Approval sets both countries side by side with a complete directory of **government departments, portals, exact statutory timelines, and verified phone numbers, faxes, emails, and street addresses** for getting official approval.
- 09. Everyday Situations: A Reader's Field Guide and 10. Primary Sources, Method, and Plain Glossary close with six everyday scenarios, a plain-language glossary, and all **26 primary sources** verified as of **2 November 2026**.

> **NOTE**
> **Note on official files:** MOFCOM published Attachments 1 and 2 in WPS format; we preserve the untouched original files in 10. Primary Sources, Method, and Plain Glossary alongside local PDF reading copies. An unrelated organoid research presentation previously saved locally as `china_export_control.pdf` was excluded from our evidence.

01. The Day Knowledge Got a Customs Booth

One Friday Afternoon, Two Very Different Clocks

Picture a city council announcing on Friday afternoon that it plans to lower the speed limit on forty stretches of highway next winter, and inviting citizens to write in with comments for thirty days. Then read the second half of the notice: starting that same Friday afternoon, a real tollbooth is already open on the highway, recording every covered truck that drives past, including trucks that left the warehouse yesterday and have not yet crossed the bridge.

That is almost literally how **Announcement No. 48 of 2026** works. Issued jointly on **16 October 2026** by China's Ministry of Commerce (MOFCOM, led by its Bureau of Industry, Security, Import and Export Control) and the Ministry of Science and Technology (MOST) under China's *Foreign Trade Law* and *Regulations on the Administration of Technology Import and Export*, the document sets two separate clocks ticking at once CN48, sections I and III CNTIER:



Fig. 6 Announcement No. 48 splits into two parallel tracks on 16 October 2026: a 30-day consultation on 40 proposed catalogue entries, and an immediate pre-export registration mandate on ecomp.mofcom.gov.cn.

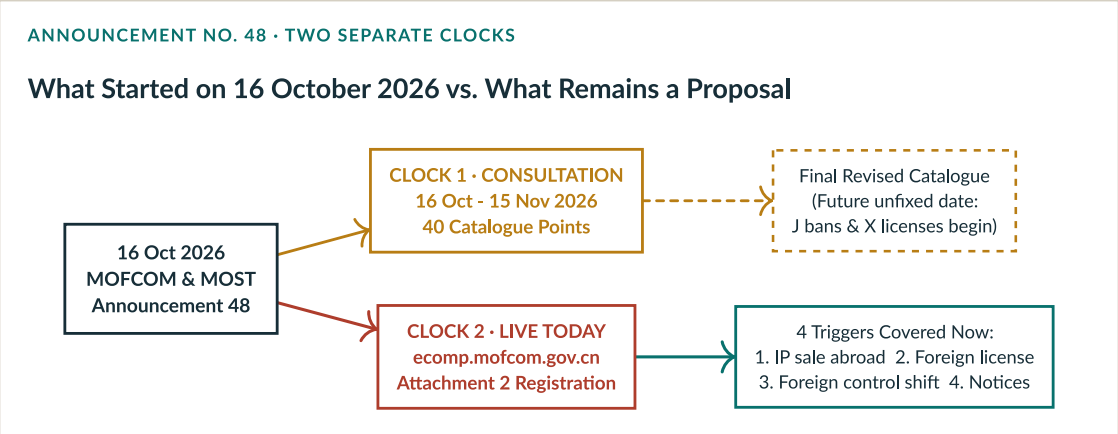


Fig. 7 Announcement No. 48 splits into two parallel tracks on 16 October 2026: a 30-day consultation on 40 proposed catalogue entries, and an immediate pre-export registration mandate on ecomp.mofcom.gov.cn.

1. **The Consultation Clock (16 October to 15 November 2026):** Beijing proposes adding **40 technology points** (14 prohibited from export and 26 restricted) across artificial intelligence and chips (23), biomedicine (5), and solar energy (12) to the *Catalogue of Technologies Prohibited or Restricted from Export*. Until the public comment window closes and a final revised catalogue is formally promulgated, those new prohibitions and licensing rules remain **proposals, not final bans** CN48, sections I and II.
2. **The Registration Clock (Running Right Now):** From the moment Announcement 48 went live on 16 October 2026 until the day the final revised catalogue eventually takes effect (a date Beijing has intentionally left unfixed), anyone exporting a technology listed in Attachment 1 must complete **pre-export information registration** online through MOFCOM's Unified Business System Platform at `https://ecomp.mofcom.gov.cn` CN48, section III(1) CNPLAT.

Confusing those two clocks is the single most common mistake people make when discussing China's new rules. Registration is active today, yet registration is a reporting gate, not an export ban. Conversely, waiting until mid-November to pay attention can catch a startup or university lab off guard if it is signing an overseas license, shipping covered equipment, or closing a foreign financing round right now CN48 CNQA.

Crossing a Border Without Packing a Single Crate

When our grandparents heard the word *export*, they pictured harbor cranes lifting steel containers onto a cargo ship in Shanghai or Rotterdam. When a modern technology regulator says *export*, not a single wooden crate needs to touch a dock.

Under Section III(2) of Announcement 48, four transactions trigger mandatory pre-export registration before completion, **including transfers that had already started before October 16 but are not yet finished** CN48, section III(2):

- **Selling the crown jewels abroad:** Transferring ownership of domestic technology to an overseas buyer, such as assigning patents or selling proprietary model or chip intellectual property.
- **Granting an overseas license:** Keeping ownership in China while licensing a foreign partner, client, or subsidiary to use covered domestic technology outside the country.
- **Changing who holds the steering wheel:** A corporate acquisition, restructuring, or cap-table shift through which the owner or *actual controller* of domestic technology changes into an overseas entity. Even if every engineer and server stays in Hangzhou or Beijing, if ultimate corporate control crosses the border, the rule engages.
- **Additional scenarios notified by MOFCOM:** A statutory catch-all allowing the ministry to address new transfer structures by official notice.

When an exporter logs into the Technology Trade Management Information Application on `ecomp.mofcom.gov.cn` to file **Attachment 2** (*Technology Export Information Registration Form*), the portal asks for concrete, auditable facts: the exporter's corporate identity and official company seal (公章), the technology owner, the foreign recipient's Chinese and English names, country, address, and corporate relationship to the exporter, the exact seven-character catalogue code and technical metrics, the end user and end use, whether the IP is self-developed, and the contract value converted into US dollars CN-A2. Multiple technologies transferred to the same overseas partner in one deal may be bundled onto one form; separate partners require separate filings CN48, section III(3).

Four Words That Change Everything

Legal catalogues look intimidating because casual readers read every line as a "ban." In practice, four regulatory categories carry vastly different meanings for human progress CN-A1, notes 1 to 4 CNQA CNTIER:

Prohibited (禁止出口, codes ending in `J`)
: A solid red wall. Once finalized, covered technology within the stated scope cannot legally be exported at all CNTIER, Art. 32.

Restricted (限制出口, codes ending in `X`)
: A guarded gate with a lock and a key. Export is lawful, provided the exporter obtains official government permission before transferring the technology CNTIER, Art. 33.

Automatic licensing (自动许可)
: Proposed facilitation for specified restricted exports. MOFCOM describes the intended process as nearly immediate, but this is distinct from an exclusion and does not itself provide a final procedure or a guaranteed approval time. CN-A1, notes 2 and 3 CNQA

Excluded (不纳入管理范围)
: Outside this catalogue altogether. Already-public open-source AI and published academic papers sit here under Footnote 1, while ordinary laws like data security or copyright remain in place CN-A1, note 1.

  Think of it like an international airport: *prohibited* items cannot fly; *restricted* items need a declared export permit; *automatic licensing* is the fast-track e-gate; and *excluded* items stay in your backpack without stopping at customs.

What This Means for the Two Teams

For our two teams, filing information before a covered transfer and being forbidden to transfer are different obstacles. The story asks what cooperation survives each rule. This chapter establishes which duty operates now and which catalogue changes remain proposed, before we judge their effects.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

02. The Map of Forty Guarded Technologies

How Governments Measure an Idea

How do you write a law that separates a student's weekend coding project from a strategic national asset? If you write "advanced AI and hardware are controlled," every calculator app and circuit board drowns in paperwork. If you name a specific company or model brand, tomorrow's rename slips straight past the gate.

Regulators solve this by drawing **numerical performance tripwires and purpose boundaries**. In Attachment 1 of Announcement 48, Beijing maps **40 distinct technology points** inside seven-character catalogue codes: **23 in artificial intelligence and AI hardware** (6 proposed prohibited, 17 proposed restricted), **5 in biomedicine** (3 prohibited, 2 restricted), and **12 in solar manufacturing** (5 prohibited, 7 restricted), matching the official totals of **14 prohibited** and **26 restricted** technologies CN48 CN-A1 CNQA.

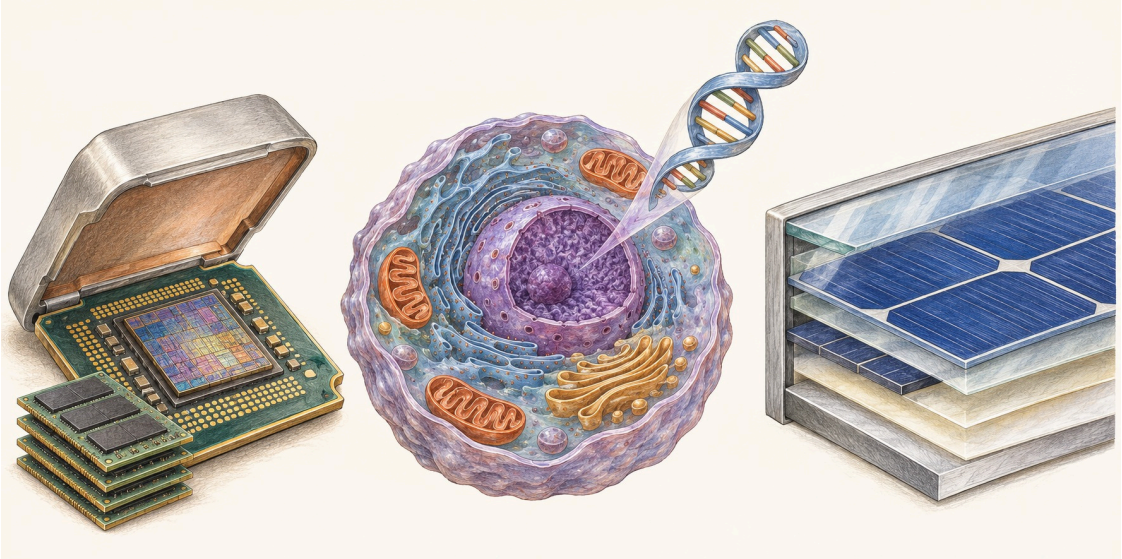


Fig. 8 Computing, biomedicine and solar manufacturing, shown through a processor, a cell and the layers of a solar panel.

Attachment 1 groups several numbered technologies under the same code. For example, 266504X has two main points, 266505X has four, 266506X has **two main points**, and 266507X has four. Our atlas splits the attachment into 40 reading checkpoints while retaining the original code and point. It is an editorial index, not a replacement legal list. CN-A1

Reading the Tripwires Without Getting Tripped

Keep three practical reading habits in mind as you explore the map:

1. **Spot "any one" versus "all together."** Entry `266504X` Point 2 (frontier large models) triggers if *any single one* of its three tests is met: $\geq 10^{26}$ FLOPs, *or* ≥ 2 trillion parameters, *or* a live Global Top-10 benchmark rank. By contrast, `266505X` Point 1 (planning agents) triggers only if an agent splits tasks into ≥ 20 steps *and* hits $\geq 80\%$ decomposition accuracy *and* reaches $\geq 90\%$ subtask completion simultaneously CN-A1.
2. **Purpose can be half the fence.** Several severe-looking entries (such as deepfake detection at $\geq 95\%$ accuracy in `266401J`, or high-performance sensitive agents in `266502J`) apply *only* when built specifically or mainly for national cross-border data defense, critical infrastructure, military, or core government settings. Ordinary studio tools and consumer assistants sit outside that purpose CN-A1.
3. **Read the hardware inside the software list.** Notice that points 18 through 23 below (`266506X` and `266507X`) are not chatbots at all: they are **processor chip architectures, 3D NAND and DRAM memory designs, and >10,000-card supercluster interconnects**, while Footnote 4 pulls **the listed categories of physical solar machinery** into export control CN-A1.

Artificial Intelligence, Chips & Superclusters (23 Checkpoints)

266401J · Point 1: Cross-border security deepfake detection *(Proposed Prohibited)*
High-grade detection systems and equipment know-how built mainly to guard China's cross-border network-data security, achieving at least 95% accuracy and 95% recall across deepfake video, audio, images, or multimodal media. Ordinary consumer media checkers outside this national-security purpose are not the target.
266402J · Point 1: Autonomous malware generation and mutation *(Proposed Prohibited)*
AI that autonomously writes or mutates malicious code capable of striking critical infrastructure, defense systems, or core government networks, hitting listed attack-success rates (at least 60% against multi-engine scanners, 85% against mainstream antivirus, or 90% against ML detectors). Everyday coding assistants are not covered.
266402J · Point 2: Self-adapting cyber-evasion AI *(Proposed Prohibited)*
Offensive AI built for sensitive national-security targets that senses sandboxes, debuggers, or emulators to escape them, or automatically rewrites at least 50% of its code structure or control flow on every run while keeping its payload intact.
266402J · Point 3: Mass automated evasion-variant swarms *(Proposed Prohibited)*
Technology for sensitive cyber operations that mass-produces and deploys evasive malware variants at scale: generating 20+ structurally distinct variants at once (or 80%+ end-to-end automation), 10+ daily variants with 30%+ syntax-tree divergence, or 90%+ behavioral consistency inside sandboxes.
266501J · Point 1: Large models trained on China's Core Data *(Proposed Prohibited)*
Large AI models trained using China's legally designated *Core Data* (the highest national-security tier under the Data Security Law) and offered as services. What triggers this red line is the legal category of training data, never a company's brand name.
266502J · Point 1: High-performance agents for critical or defense settings *(Proposed Prohibited)*
AI agents built specifically or mainly for critical infrastructure, national defense, or core government operations that also reach the high-capability benchmarks defined in `266505X`. General-purpose productivity or coding agents do not fall here.
266402X · Point 8: Ultra-high-accuracy deepfake detection (98%+) *(Proposed Restricted)*

Deepfake detection systems and equipment-manufacturing know-how for video, audio, images, or multimodal content reaching both accuracy and recall of at least 98%, regardless of whether used for border-data defense.

266402X · Point 9: Near-invisible video and image steganography *(Proposed Restricted)*
Hiding and extracting secret data inside images or video only when all three benchmarks are met together: high survival under distortion (bit-error <=0.5% or recall >=99.9%), false-alarm rate <=0.1%, and near-total invisibility to human eyes (SSIM >=0.999, PSNR >=41 dB, or LPIPS <=0.005).

266404X · Point 6: AI-assisted zero-day and exploit discovery *(Proposed Restricted)*
Cybersecurity AI meeting any of six high bars: >=75% on top exploit benchmarks; building full exploit chains (>=50% on major OS/browsers or >=40% arbitrary code execution); 10x speed over senior human researchers; chaining 2+ low/medium bugs into working PoC attacks; finding 5+ critical or 50+ new bugs in a major OS or open-source project in one pass; or reaching the top tier of mainstream exploit scales.

266501X · Point 19: High-speed, anti-spoof or dialect voiceprint ID *(Proposed Restricted)*
Speaker-identification algorithms and systems hitting any one of four bars: >=99.1% true-match rate at <=0.01% false positives on short clips; <=300 ms latency at >=100 requests/sec; >=98% spoof detection (or <=1% false acceptance) against replay and synthetic voices; or >=90% accuracy tailored to Chinese minority languages or regional dialects.

266501X · Point 20: Chinese and minority-language machine translation *(Proposed Restricted)*
Translation technology built mainly for Chinese and China's ethnic-minority languages that achieves either a human quality score (MOS) of at least 4.8, or >=95% semantic fidelity paired with BLEU >=35 in general text (or BLEU >=30 in specialist domains).

266504X · Point 1: Large models trained on China's Important Data *(Proposed Restricted)*
Large models trained on legally defined *Important Data* and offered as services. Crucially, Note 1 carves out trace use (<=0.1%) when sanitized so the data neither boosts core domain R&D capability nor reappears in outputs, as well as data lawfully obtainable from public sources.

266504X · Point 2: Frontier large-model core technology *(Proposed Restricted)*
The heart of the AI model rule: covers weights, key checkpoints, underlying source code, architecture, training optimization, and inference acceleration if ANY ONE test is met: >=10^26 training FLOPs, >=2 trillion parameters, or a Global Top-10 public ranking at export time. Open-source releases qualify for Footnote 2 automatic-licensing facilitation.

266505X · Point 1: Long-horizon autonomous planning agents *(Proposed Restricted)*
Task-planning technology that breaks natural-language instructions into chains of 20 or more executable subtasks while hitting at least 80% decomposition accuracy and at least 90% average subtask completion. All three numbers must be met together.

266505X · Point 2: 60-agent collaborative orchestration *(Proposed Restricted)*
Multi-agent coordination frameworks running 60 or more concurrent AI agents at once, finishing work in no more than twice the fastest single-agent time for equivalent quality, and reaching >=85% agreement with ground-truth answers.

266505X · Point 3: 200-tool operating agents *(Proposed Restricted)*
Tool-use and system-operation technology supporting 200+ integrated third-party tools, picking the right tool at least 90% of the time, and succeeding at least 90% end-to-end on complex tasks that chain three or more distinct tools.

266505X · Point 4: Real-time strategic game and decision agents *(Proposed Restricted)*
Strategic AI meeting all three criteria at once: >=20 planning steps (or >=10 concurrent agents); adaptive sub-200 ms decisions across >=10 control units in dynamic environments; and outperforming 90% of human players on standard strategic benchmarks.

266506X · Point 1: AI processor architecture, instruction sets & IP cores *(Proposed Restricted)*
AI processor design technology covering custom architectures, instruction sets, compute IP cores, and base software stacks meeting any listed trait, including <=14 nm process need, floating-point

or integer compute capability, 2.5D/3D chiplet packaging, high-speed interconnects, non-von-Neumann / compute-in-memory designs, or named vision/voice/driving uses. Footnote 2 offers automatic-licensing facilitation for overseas foundry tape-outs.

266506X · Point 2: Advanced 3D NAND flash and stacked DRAM / HBM chip design *(Proposed Restricted)*
3D NAND flash memory design hitting >=8 Gb/mm2 density, >=1.5 Gbps I/O speed, or wafer bonding; or DRAM memory design using <=18 nm lithography, vertical-channel transistors, bonding, or multi-layer HBM-style wafer/die stacking.

266507X · Point 1: Ultra-fast 10,000-GPU cluster interconnects *(Proposed Restricted)*
Chip-to-chip or node-to-node interconnect technology delivering bidirectional bandwidth of at least 800 GB/s inside giant AI compute clusters of more than 10,000 accelerator cards.

266507X · Point 2: Rapid fault recovery in 10,000-card training *(Proposed Restricted)*
Distributed-training resilience in >10,000-card clusters meeting any bar: checkpoint recovery (MTTR) <=120 seconds, seamless task migration with <=30 seconds of visible interruption, or <=5% effective GPU compute lost to network faults.

266507X · Point 3: Microsecond cluster communication scheduling *(Proposed Restricted)*
Communication scheduling for >10,000-card AI clusters achieving either >=98% linear scaling, dynamic topology-aware routing with <=1 microsecond end-to-end latency, or >=90% overlap hiding communication behind computation.

266507X · Point 4: Mixed-vendor and cross-region GPU orchestration *(Proposed Restricted)*
Software and ops know-how uniting AI chips from two or more different hardware vendors inside >10,000-card clusters, hitting listed high-bandwidth, throughput (+30%), card-utilization (+65%), cost-drop (-20%), 99.9% uptime, or >=80% cross-data-center training efficiency marks.

Biomedicine & Gene Platforms (5 Checkpoints)

262704J · Point 1: Reusable cell and gene therapy (CGT) platforms *(Proposed Prohibited)*
Chinese-owned foundational platforms capable of generating multiple cell or gene therapies. Crucially, only selling ownership abroad is prohibited: out-licensing, CRO/CDMO contract manufacturing, and tech needed to develop, test, or market a specific individual medicine remain outside the ban.

262705J · Point 1: Reusable siRNA drug-making platforms *(Proposed Prohibited)*
Chinese-owned foundational small-interfering RNA (siRNA) platforms, governed by the exact same boundary: transferring ownership of the reusable engine abroad is barred, while licensing deals, contract R&D/production, and single-drug partnerships stay open.

267302J · Point 1: Pathogen enhancement and directed evolution *(Proposed Prohibited)*
Biosecurity red line covering artificial genetic modification or controlled serial-passage techniques that enhance pathogen function, sustain virulence at scale, or accelerate evolutionary adaptation. Routine medical diagnostics and vaccine care are unaffected.

262705X · Point 1: Discovery of unseen cell and gene therapy targets *(Proposed Restricted)*
Chinese-owned methods for finding and validating wholly new therapeutic targets never previously disclosed anywhere in the world via patents, journals, public databases, or conferences. Footnote 3 offers automatic-licensing facilitation for overseas licensing, joint R&D, and public-health emergencies.

262706X · Point 1: Discovery of unseen siRNA drug targets *(Proposed Restricted)*
Chinese-owned methods for screening and locking in wholly new, globally undisclosed siRNA drug targets, paired with the same Footnote 3 fast-track facilitation for international licensing, collaborative research, and humanitarian health responses.

Solar Silicon, Wafers & Factory Machinery (12 Checkpoints)

263003J · Point 1: Mega-scale cold-hydrogenation polysilicon lines *(Proposed Prohibited)*
Cold-hydrogenation fluidized-bed recipes and equipment-making know-how capable of producing at least 250,000 tonnes of trichlorosilane per single line each year.

263003J · Point 2: High-capacity silane fluidized-bed processes *(Proposed Prohibited)*
Silane fluidized-bed chemical recipes and equipment-making technology with single-line output of at least 20,000 tonnes per year.

263004J · Point 1: Paper-thin 70-micrometer silicon wafers *(Proposed Prohibited)*
Know-how and machine-building technology for slicing solar silicon wafers 70 micrometers thick or thinner (roughly the width of a human hair), explicitly including control software, underlying algorithms, and engineering drawings.

263801J · Point 1: Edge passivation for thin large-format cells *(Proposed Prohibited)*
Solar-dedicated edge-passivation processes and toolmaking know-how when used on 182 to 230 mm crystalline-silicon wafers that are 140 micrometers thick or thinner.

263801J · Point 2: 30%-efficient perovskite and tandem solar tech *(Proposed Prohibited)*
Next-generation solar manufacturing across four benchmarks: >=30%-efficient mass-produced silicon-perovskite tandem cells (>=182 mm); >=1.1 m wide slot-die coating; >=18%-efficient single-junction perovskite modules (>=1 m2); or >=28%-efficient tandem modules (>=1 m2).

263004X · Point 1: 30,000-tonne polysilicon plant engineering *(Proposed Restricted)*
Complete plant-engineering packages, factory blueprints, and equipment designs for single-line polysilicon lines producing 30,000 tonnes or more per year.

263004X · Point 2: 800-tonne-per-furnace reduction processes *(Proposed Restricted)*
High-yield reduction furnace recipes and equipment-manufacturing technology turning out at least 800 tonnes of polysilicon per single furnace each year.

263005X · Point 1: Large-chamber monocrystalline silicon pullers *(Proposed Restricted)*
Solar ingot crystal-growing furnace designs hitting any listed mark: chamber inner diameter >=1,600 mm, thermal field >=35 inches, pull speed >=1.2 mm/min, or equal-diameter power draw <=40 kW.

263005X · Point 2: Long silicon-ingot cutting and grinding *(Proposed Restricted)*
Half-ingot cropping/grinding processes or precision machining and toolmaking know-how for solar silicon ingots 720 mm or longer.

263005X · Point 3: Rectangular or 130-micrometer thin-wafer slicing *(Proposed Restricted)*
Manufacturing know-how for rectangular solar wafers (>=182 mm on any side) or thin wafers (<=130 micrometers) using <=28-micrometer tungsten wire saws or high-speed spindles (>=6,500 rpm) and compact roller geometry.

263802X · Point 1: High-efficiency TOPCon and back-contact (XBC) cells *(Proposed Restricted)*
Six industrial process families (LPCVD/PECVD tunnel oxide, fast ALD passivation, boron diffusion, copper plating, microcrystalline deposition, and XBC plating) on 182 to 230 mm wafers <=140 micrometers thick for >=25%-efficient cells or >=22.3% modules.

263802X · Point 2: High-efficiency heterojunction (HJT) cell processes *(Proposed Restricted)*
Four HJT manufacturing families covering low-indium (<=50%) or ultra-fast TCO coating, crystalline-silicon film deposition on <=110-micrometer wafers, low-silver (<=30%) silver-coated copper pastes, and copper-plated HJT electrodes.

What This Means for the Two Teams

The atlas counts proposed technology points. The game scores express invented preferences; they do not count technologies. For the two teams, the useful question is which precise entry and conditions cover the next contribution, rather than whether all AI sits behind one wall.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

03. Mind and Silicon: Guarding Models, Chips, and Clusters

The Sourdough Starter, the Bakery Window, and the Cast-Iron Oven

To an everyday person typing a question into a phone, artificial intelligence feels like one seamless magic trick: a text box that talks back. To an engineer, and now to export-control writers in Beijing and Washington, an AI system has **three distinct physical and digital layers**:

1. **The bakery window (Remote API & apps):** Handing a customer a warm slice of bread through a shop window tells them nothing about how to wire the oven or bake ten thousand loaves of their own tomorrow.
2. **The sourdough starter and master recipe (Weights, checkpoints & training code):** Billions or trillions of learned numerical parameters (模型权重), intermediate training snapshots (关键检查点), model architecture, and training/inference optimization code (训练优化、推理加速). Copy these files, and you replicate the trained mind itself.
3. **The cast-iron oven and power plant (AI chips, memory stacks & superclusters):** The physical silicon processors, high-bandwidth memory, chip blueprints, and optical cluster networks without which neither training nor private local deployment can happen.

MOFCOM's Q&A distinguishes remote API, app, web, SaaS and lightweight local services from transfers of restricted large-model technology under the proposed measures. This is a statement about that scope, not an exemption from every technology or data rule. Attachment 1 separately lists model technology (`266504X`), processor and memory designs (`266506X`), and specified technologies for clusters with more than 10,000 accelerator cards (`266507X`). CNQA CN-A1

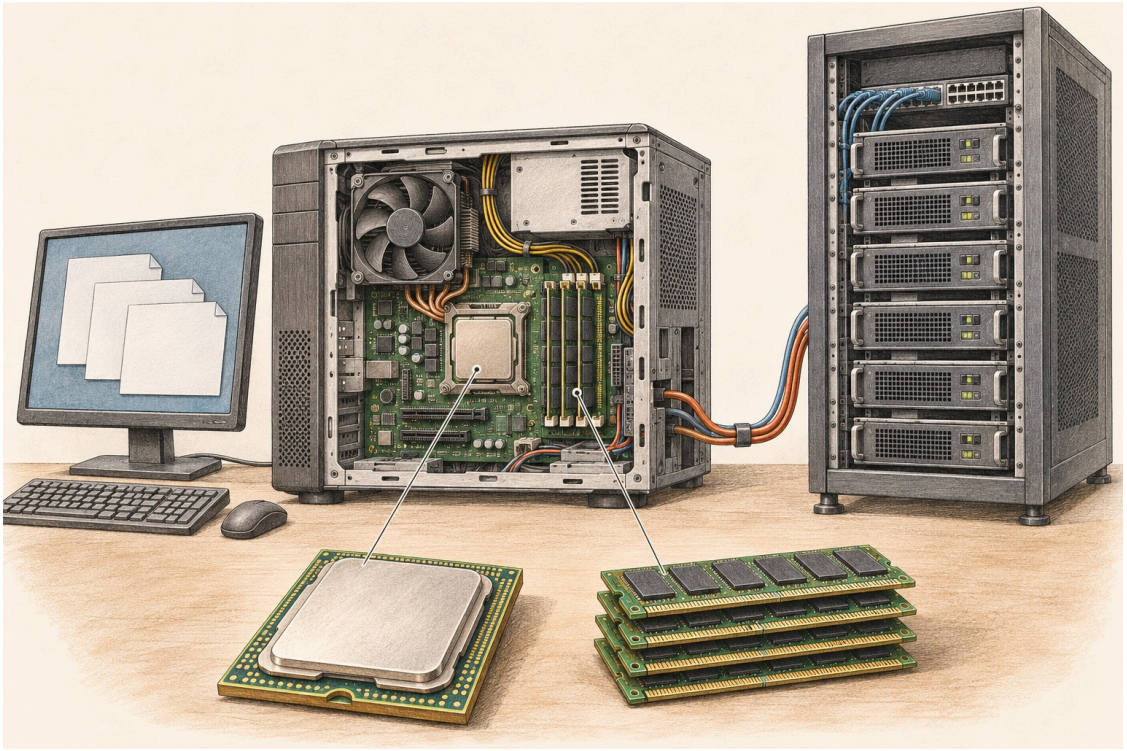


Fig. 9 Open models depend on chips, memory and computing infrastructure. A restriction on one part can affect access to the whole system.

Three Tripwires for a Frontier Model (`266504X`)

Not every neural network trained in a university lab or startup garage crosses into proposed restricted entry **266504X Point 2**. A model's core technology (weights, key checkpoints, underlying source code, architecture, training optimization, and inference acceleration) enters scope if it crosses **any one** of three tripwires CN-A1, 266504X(2), note 3:

1. **Compute scale:** Total training computation reaches at least 10^{26} **floating-point operations (FLOPs)**, the massive compute threshold also used in America's `4E091` model-weight rule US25.
2. **Parameter size:** Total parameter count reaches at least **2 trillion (2,000 billion)**.
3. **Global Top-10 capability at the moment of export:** Even a lean, compact model qualifies if, **at the time export occurs**, public evaluations by mainstream domestic or international benchmark institutions or major industry platforms rank it in the **Global Top 10** across overall capability, specialist business tasks, global call volume, safety performance, or training efficiency CN-A1, 266504X, note 3.

That third tripwire is uniquely alive. Because open-source researchers constantly invent clever ways to distil frontier reasoning into 30B or 70B parameters, a compact prodigy can step *into* restricted scope on the week it debuts at #6 in the world, and step *back out* eight months later when newer models overtake it. Classification belongs to a specific model version on a specific date, never to a brand name like "DeepSeek" or "Qwen" forever CN-A1.

Data Red Lines (`266501J` & `266504X(1)`) and Autonomous Agents (`266505X` & `266502J`)

Even a small model hits a hard wall if trained on sovereign secrets. Borrowing categories from China's *Data Security Law*, Attachment 1 places models trained on legally designated **Core Data**

behind a **proposed prohibition** (`266501J`), even for remote service provision, while placing models trained on **Important Data** in the **restricted lane** (`266504X(1)`), with explicit carve-outs when the Important Data came from **lawful public channels** or formed **at most a 0.1% sanitized trace share** incapable of boosting domain R&D or leaking into outputs CN-A1.
Surrounding the base model, four restricted points in **266505X** cover high-end **AI agents**: task planners ($\geq$ 20 steps at $\geq$ 80% decomposition accuracy and $\geq$ 90% completion), multi-agent swarms ($\geq$ 60 concurrent agents), tool-calling engines ($\geq$ 200 external software tools at $\geq$ 90% multi-tool accuracy), and sub-200-millisecond decision systems beating 90% of human operators. Only when built *specifically or mainly* for defense, critical infrastructure, or core government settings do such agents cross into prohibition (`266502J`) CN-A1.

The Hardware Layer Inside China's AI List: Chips, Memory, and 10,000-Card Clusters

The AI section also covers processor designs, memory designs and large-cluster engineering:

- **AI Processor Blueprints (266506X Point 1):** Proposes restricting the export of AI processor architectures, custom instruction sets, IP cores, base software stacks, and design tools meeting *any* listed trait: not only $\leq$ 14 nm process designs, 2.5D/3D chiplet packaging, or novel non-von-Neumann / compute-in-memory architectures, but even broad functional traits such as floating-point/integer compute or smart-voice, vision, and autonomous-driving uses CN-A1, 266506X. Think about what that touches: open-source hardware initiatives (such as RISC-V AI accelerators), custom NPU startups, and cross-border chip engineering teams! Because fabless chip teams must send finished layouts to overseas semiconductor foundries to manufacture physical silicon, **Footnote 2** had to build a vital escape valve: **automatic-licensing facilitation for overseas chip tape-out** (赴境外开展芯片流片制造) CN-A1, note 2 CNQA.
- **Memory designs (266506X, Point 2):** The point contains separate subpoints for 3D flash and memory-chip design. Flash tests include density, I/O speed or bonding; memory tests include a process of 18 nm or smaller, vertical channels, bonding or multilayer stacking. These are alternative characteristics, not a single 18 nm threshold for every kind of memory. CN-A1, 266506X(2)
- **Supercluster Plumbing (266507X Points 1 to 4):** Training a frontier open model on **more than 10,000 accelerator cards** requires industrial wizardry so the cluster does not stall every hour. Four restricted points guard $\geq$ 800 GB/s node interconnects, $\leq$ 120 -second automated fault recovery, $\leq$ 1 -microsecond collective routing, and **heterogeneous cluster software that combines chips from two or more hardware vendors** into one training pool CN-A1, 266507X. When Washington restricts Nvidia GPUs (`3A090`) and Beijing restricts multi-vendor cluster stitching (`266507X`), open-source distributed-training frameworks sit right on the fault line between both regimes!

What This Means for the Two Teams

A's memory-saving improvement and B's tests are knowledge the teams might exchange. Their computers are a separate dependency. An open-model license cannot supply a missing chip, and a hardware restriction does not by itself describe the legal status of every model file. We follow both parts of the workbench.
The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

04. Open Models: Public Releases and Private Transfers

An Open License Is One Part of a Release

A team can publish code and weights for others to inspect, adapt and run. A permissive copyright license is one part of that release. Applicable export, data and safety rules still need assessment. DeepSeek-R1 and Qwen3 are documented examples of public releases; their publication does not establish an exemption for every related activity. R1 QWEN



Fig. 10 Receiving files and knowledge to work independently differs from using a service that keeps its underlying model elsewhere.

Start with the Actual Artifact

A public paper, downloadable weights, private training code and a later fine-tune are different artifacts. The Open Source AI Definition asks for freedoms and access to code, parameters and sufficient data information. A product called open source may provide only some of these. A copyright label alone also does not determine export-law status. OSI

China: Registration Now, Catalogue Proposal Separately

Announcement 48 requires temporary registration before covered transfers, including unfinished exports. It identifies ownership transfers, overseas licensing, foreign ownership or control changes and further notified cases. The proposed catalogue prohibitions and licensing categories are a separate track. A proposed prohibited entry is not, by itself, a new ban already in force. Existing law still needs to be checked. CN48 CN-A1

We read the Chinese draft with its section-wide notes, rather than treating every release as a compulsory sequence of four gates:

1. **Already-public AI technology:** Note 1 excludes technology already public through open source or papers. Identify exactly what was published. A paper about private weights does not publish those weights; a private derivative does not automatically inherit every exclusion. CN-A1, note 1
2. **Unpublished data-sensitive technology:** The draft places specified core-data models in a proposed prohibited category and important-data models in a restricted category. The trace-data exclusion has several conditions, not just a 0.1% test. These categories need exact facts and legal interpretation. CN-A1, 266501J and 266504X
3. **Other unpublished controlled technology:** Check alternative model thresholds and any agent, chip-design or cluster technology being transferred. Being below one model threshold does not clear every entry, existing law or bundled artifact. CN-A1
4. **Facilitated routes:** Note 2 proposes automatic licensing or similar facilitation for specified restricted AI exports, including open-source releases, overseas chip fabrication and certain self-operated computing centers. An exclusion and an easier permission route are different. Final procedures and applicable conditions still matter. CN-A1, note 2 CNQA

Table 2 Three separate questions before a transfer

Question	What we check
Current Chinese registration	The covered technology and transfer, including unfinished exports.
Proposed Chinese catalogue	The precise entry, its notes, exclusions and proposed facilitated routes.
Separate US rules	Whether the artifact is subject to the EAR, followed by its classification and applicable conditions.

United States: A Separate Jurisdiction Test

First establish whether the artifact is subject to the EAR. Published-information provisions and the 4E091 notes can exclude qualifying public technology or weights. The model-weight entry has additional-training conditions for derivatives, so a privately improved model is not automatically treated like its public base. The announced non-enforcement policy and GAO's decision must also be distinguished from the published text. EAR734 EAR774, 4E091 notes GAO26

MOFCOM's answer on API, app, web, SaaS and lightweight local deployment concerns restricted large models under the proposed measures. It should not be turned into a universal US exemption or a clearance for other Chinese controls. Handing private weights to another party is also different from providing a remote service. CNQA

Possible Effects on Open Models

Continued public releases are consistent with the draft's explicit facilitation. **More documentation or delay** is a possible result of classification work. **A public smaller model alongside a private service** is a possible business choice. **Greater concentration of training** is a possible

result if affordable computing becomes harder to obtain. These are mechanisms to investigate, not outcomes established by the announcement or measured forecasts. CN-A1 CNQA NTIA

We should not claim that chip restrictions caused a particular model's efficiency improvements without evidence isolating that cause. Smaller parameter counts do not guarantee exclusion because other criteria may apply. Local use can reduce reliance on a remote service, but it does not by itself guarantee privacy, security or clinical reliability.

Independent Review and Its Limits

Open access can let teams outside the original developer inspect a model and suggest repairs. More reviewers can help, but they can share the same blind spots or lack resources to test well. Finding a flaw is also different from fixing it. The optional probability example in Optional: Independent Review Is Not a Safety Guarantee explains these limits.

Open weights can also help harmful actors modify a system or remove safeguards. NTIA's 2024 report considered both benefits and risks, recommending monitoring rather than restricting then-current open weights. Research on open foundation models stresses the need to compare the extra risk with existing alternatives. Neither source establishes a universal safety multiplier for openness. NTIA OPENRISK

A Practical Release Record

Our suggested record identifies the model version, file hashes, license, public and private artifacts, data categories, measured criteria, recipient, destination and transfer form. Keep publication evidence and any relevant filing or permission. This is a documentation suggestion, not a new compulsory form or a substitute for the official registration fields. CN-A2

What This Means for the Two Teams

Our teams keep their public starting model. Their next contribution could be a paper, new code, private weights or a service, each needing its own assessment. Public-information exclusions and proposed facilitated routes matter because they may preserve specific exchanges that the story's broad wall would stop.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

05. Living Code: Gene Platforms and Shared Cures

Guarding the Printing Press While Sharing the Books

Artificial intelligence writes in floating-point numbers; modern biomedicine writes in the four-letter alphabet of life (`A, C, G, U/T`). Over the past decade, Chinese biotech teams have advanced from making off-patent generics to inventing reusable **cell and gene therapy (CGT)** engines and **small-interfering RNA (siRNA)** platforms that silence disease-causing genes before harmful proteins form.

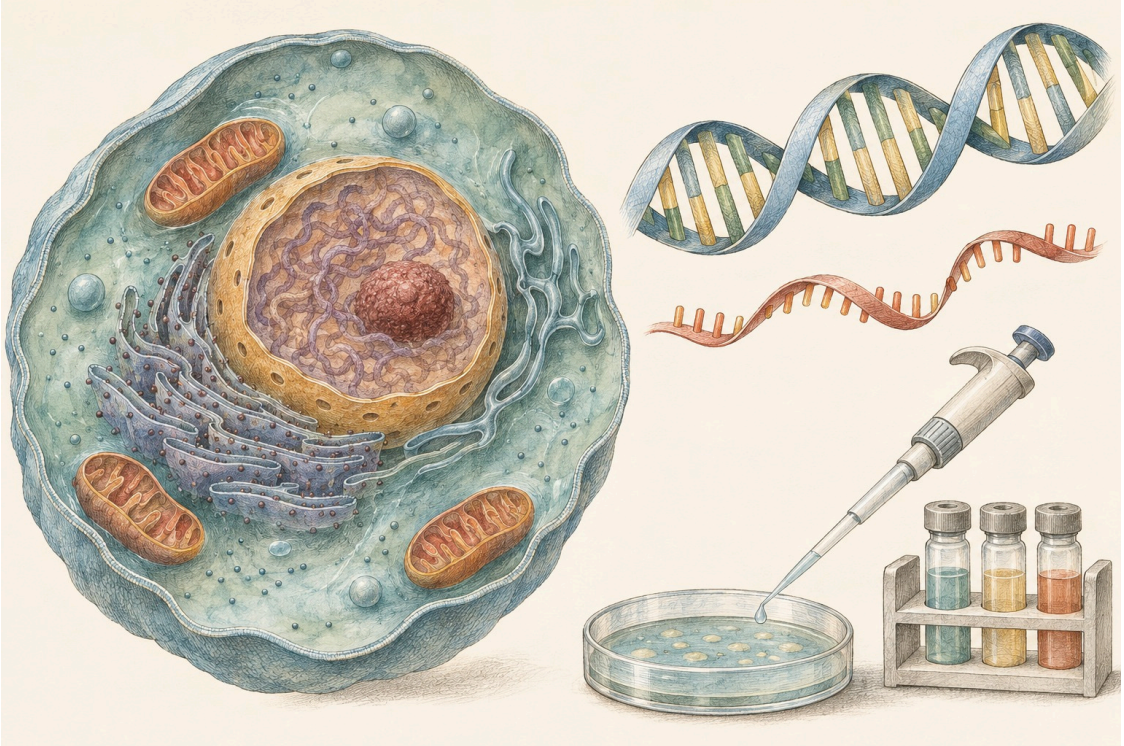


Fig. 11 A cell, DNA, RNA and laboratory tools make the subject tangible. A reusable medicine-development platform differs from a particular treatment.

The draft distinguishes ownership of a reusable medicine-development platform from specific drug projects. It proposes limits on covered platform ownership transfers, with explicit exclusions for licensing and other activities. Those exclusions concern these entries; medicines and clinical trials can still face other rules. CN-A1, 262704J and 262705J CNQA

Picture a foundational CGT or siRNA platform as a custom **printing press** perfected over ten years (delivery vectors, chemical stabilization recipes, and modular design rules capable of printing twenty future therapies). Picture a specific clinical drug candidate as **one book** printed on that press:

- **Platform ownership transfer is proposed prohibited (262704J and 262705J).** The draft targets overseas ownership transfers of the specified reusable platforms, with explicit exclusions described below. The new prohibition is a proposal, not an already-final rule introduced by Announcement 48. CN-A1
- **Out-licensing, contract manufacturing, and single-drug partnerships stay open:** Three explicit notes under both entries carve out overseas **out-licensing** (`对外授权许可`), **outsourced R&D and manufacturing (CRO/CDMO)** where ownership remains domestic, and all pipeline data needed for **preclinical research, clinical trials, manufacturing, quality control, regulatory approval, and commercialization of specific drug candidates** in global business-development deals CN-A1 CNQA. Only a sham "perpetual buyout license" that strips domestic ownership in substance crosses the red line.

Undisclosed Biological Targets and Biosecurity Walls

Two **proposed restricted** entries (`262705X` and `262706X`) cover Chinese-owned technology for discovering and validating **wholly new therapeutic targets** in CGT and siRNA medicine, defined strictly as molecular targets **never previously disclosed anywhere in the world** in patents, journals, public databases, industry filings, or conferences CN-A1. Even there, **Footnote 3** provides **automatic-licensing facilitation** for public-health emergencies, humanitarian relief, international collaborative research using overseas resources, and overseas out-licensing CN-A1, note 3 CNQA.

The fifth entry (`267302J`, proposed prohibited) is a pure **biosecurity wall** against artificial genetic modification and directed serial-passage methods that enhance pathogen virulence or accelerate controlled pathogen evolution CN-A1, 267302J, paralleling Category 1 pathogen controls and `EAR 744.4` on the US side EAR774 EAR744. *(Bulk human genomic data, meanwhile, faces separate data-security rules in both countries, including the US Justice Department's Data Security Program `DOJ <https://www.justice.gov/nsd/data-security>`_.)*

What This Means for the Two Teams

In medicine, the shared workbench involves platforms, research and particular treatments. The distinctions above help us ask which collaboration a rule reaches and which it leaves possible. We do not transfer the AI game scores to medicine or assume that open access alone establishes patient safety.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

06. Capturing Sunlight Inside a Machine

You Can Ship the Solar Panel, Not the Secret Kitchen

The official explanation distinguishes manufacturing technology from products made with it. Announcement 48 does not create a blanket ban on finished solar panels. That statement concerns this announcement, not every customs rule or other measure that could affect a shipment. CNQA



Fig. 12 A crystal-growing furnace, thin wafers and a coating machine show the equipment and knowledge behind a finished solar panel. The machinery is schematic.

Why, then, do **12 of the 40 technology points** (5 proposed prohibited, 7 proposed restricted) target solar manufacturing? Because China leads the world in turning raw quartz into ultra-pure polysilicon, pulling giant monocrystalline ingots, slicing wafers thinner than office paper, and coating next-generation tandem cells. As foreign governments offer subsidies and erect import tariffs to build domestic solar factories, Beijing is drawing a border around **both the industrial recipes and the physical factory machines that carry them** CN-A1 CNQA:

1. **Silicon purification (263003J & 263004X):** Proposed prohibitions cover mega-scale cold-hydrogenation fluidized beds ($\geq 250,000$ tonnes/year trichlorosilane on a single line) and silane fluidized beds ($\geq 20,000$ tonnes/year); restricted entries cover $\geq 30,000$ -tonne plant packages and ≥ 800 -tonne-per-furnace reduction recipes CN-A1.

2. **Ingot pulling and paper-thin wafers (263004J & 263005X):** Slicing silicon wafers $\leq$ 70 micrometers thick (including control software, algorithms, and mechanical drawings) is **proposed prohibited** (`263004J`). Restricted entries cover giant crystal pullers ($\geq 1,600$ mm chambers or ≤ 40 kW power draw), machining ≥ 720 mm ingots, rectangular wafers ($\geq$ 182 mm), and ≤ 130 -micrometer slicing with ≤ 28 -micrometer tungsten wire saws CN-A1.
3. **Tandem and high-efficiency cells (263801J & 263802X):** Proposed prohibitions guard mass-produced **perovskite and silicon-perovskite tandem** recipes ($\geq 30\%$ cell efficiency, ≥ 1.1 -meter slot-die coaters, $\geq 28\%$ large tandem modules) and thin-wafer edge passivation. Restricted entries cover high-throughput **TOPCon, back-contact (XBC), and heterojunction (HJT)** processes ($\geq 25\%$ cell or $\geq 22.3\%$ module efficiency) CN-A1.

Footnote 4: When the Physical Machine *Is* the Recipe

Just as AI software cannot be separated from AI chips, solar know-how cannot be separated from solar furnaces. When a supplier ships a deposition reactor, installs embedded control code, and trains foreign engineers for six months on the factory floor, the hardware carries the mind inside it. **Draft Note 4** covers specified core photovoltaic equipment when it embodies listed technical indicators, including reactors, crystal furnaces, wafer saws and coating machines. Covered exports require temporary pre-export registration now. If adopted, the final catalogue would distinguish prohibited technology from restricted technology that needs permission. A license would not make a prohibited transfer permissible; the final text and procedures must be checked. CN-A1, note 4 CNQA CNMEAS

What This Means for the Two Teams

A finished panel, a manufacturing machine and the knowledge used to operate it are different things. For a cooperating factory team, access to one does not supply all the others. The workbench question becomes which equipment and know-how can still move together under the applicable rules.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

07. Washington's Gatekeepers: Silicon Chips and Model Weights

Guarding the Physical Engines of Global Computing

Cross the Pacific to Washington, D.C., and you meet the older, hardware-first half of the global border architecture. Administered by the Department of Commerce's **Bureau of Industry and Security (BIS)** under the **Export Administration Regulations (15 CFR Parts 730-774)**, the US system governs global access to **frontier computing hardware, chip-design software, and semiconductor manufacturing tools** EAR734 EAR742 EAR774.
Every project touching the US system faces four sequential questions:

1. **Is it subject to the EAR?** Part 734 covers items in the US, US-origin items abroad and specified foreign-made items, subject to stated exclusions and conditions. The published-information and fundamental-research provisions have qualifications, including special treatment for encryption. A foreign chip is not covered merely because some US software was involved. EAR734
2. **What is its ECCN (Part 774)?** Controlled items carry codes on the Commerce Control List, notably **3A090** for advanced AI accelerators and **4E091** for certain closed frontier AI model weights EAR774.
3. **Where is it going (Part 742)?** Licensing turns on destination country groups.
4. **Who is the end user, and what is the end use (Part 744)?** Shipments to parties on the **Entity List (744.11)** or supporting restricted supercomputer, advanced-node semiconductor, or weapons end uses (`744.23`, `744.4`) require licenses even when an item would otherwise be unrestricted EAR744.

The January 2026 Semiconductor Rule: Permission to Knock, Not an Open Door

For three years, high-end AI data-center chips bound for China faced presumptive denial. On **15 January 2026**, BIS issued a targeted hardware calibration (`91 FR 1684`, revising `15 CFR 742.6(b)(10)`): qualifying exports from the United States to end users in China or Macau may apply for **case-by-case license review** rather than automatic denial if the chip sits below **Total Processing Performance (TPP) of 21,000** and **total DRAM bandwidth of 6,500 GB/s**, a performance window exemplified in the rule by Nvidia's **H200** and AMD's **MI325X** US26 EAR742.
Yet read the conditions carefully: case-by-case review is permission to apply, not a free pass. Applicants must certify that US domestic customers face no supply shortage, that global foundry

capacity is not diverted away from US buyers, that aggregate shipments stay within strict ratios of US volume, that rigorous physical and cybersecurity KYC controls block unauthorized remote or Entity-List use, and that an **independent US-headquartered laboratory** tests and verifies the chips' performance before export US26, sections I and IV. Third-country reexports do not automatically enjoy the same review lane.

The Strange Legal Life of America's AI Model-Weight Rule (`4E091`)

When Washington tried to move from hardware chips up into **AI model weights**, it produced a remarkable three-act legal drama:

- **Act 1 (15 January 2025):** BIS published the **Framework for Artificial Intelligence Diffusion** (`90 FR 4544`), creating ECCN **4E091** to license exports of closed model weights trained with at least 10^{26} computational operations, while **expressly excluding publicly available open-weight models** (and models below the leading open-weight benchmark) US25 EAR774, 4E091.
- **Act 2 (13 May 2025):** BIS announced it was **directing enforcement officials not to enforce** the AI Diffusion Rule pending a formal replacement rulemaking USNE.
- **Act 3 (12 May 2026):** GAO Decision `B-337935` concluded that the announced non-enforcement policy was a rule requiring submission under the Congressional Review Act. GAO separately recorded Commerce's explanation that rescission had been initiated, not completed. This was a decision about congressional review, not a court judgment repealing or enforcing `4E091`. We separately checked that the entry remained in BIS's published Part 774 text on 2 November 2026. GAO26 EAR774

The public-weights provisions are important for open-model developers. They do not mean every modification of a public model automatically remains excluded: 4E091 has specific derivative-model notes. Hardware access, software or data restrictions and the distinction between the rule text and announced non-enforcement remain separate questions. EAR734 EAR774 GAO26

Two Neighboring Washington Gates: Venture Dollars and Bulk Sensitive Data

Beside BIS stand two separate federal gatekeepers:

- **US Treasury Outbound Investment Security Program:** Since **2 January 2025**, US persons funding Chinese semiconductors, quantum tech, or certain AI systems face prohibition or mandatory notification; on **22 October 2026**, Treasury opened public comments (through **7 December 2026**) on proposed refinements to technical thresholds TREAS.
- **US Department of Justice Data Security Program:** Restricts bulk cross-border transfers or access involving sensitive US personal data (genomic, biometric, health, financial, geolocation) and government-related data DOJ.

What This Means for the Two Teams

Country A and Country B in the game are fictional. US and Chinese controls have different scopes, institutions and positions in the computing supply chain. This chapter explains those differences; it does not assign either real country a score of 3, 5, 17 or 20.
The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

08. Two Border Systems and How to Get Official Approval

Two Philosophies Meeting Inside the Same Project

Set China's Announcement 48 beside America's Export Administration Regulations and you see two superpowers guarding opposite ends of the same bridge. Beijing guards **outward transfers of domestic ingenuity and corporate ownership**, asking whether Chinese-developed model cores, chip blueprints, gene platforms, or solar machines are leaving China or passing under foreign control. Washington guards **global access to physical computing bottlenecks and restricted end users**, tracking chips, tools, and Entity-List parties worldwide CN48 EAR734 EAR742 EAR744.

Table 3 Side-by-side comparison of China's Announcement 48 and US BIS EAR controls

Dimension	China (MOFCOM & MOST Announcement 48)	United States (BIS Export Administration Regulations)
Core legal hook	Outward transfer of domestic technology ownership, overseas licensing, or change of actual controller into a foreign entity CN48 CNTIER.	US-origin items, items in the US, and foreign-made items built with US chip software/tools under Foreign Direct Product Rules EAR734.
Hardware & silicon	Proposed `266506X` (<=14 nm AI chips, chiplets, <=18 nm DRAM/HBM, 3D NAND), `266507X` (>10,000-card clusters), and listed solar equipment categories CN-A1.	ECCN `3A090` AI chips, fab equipment, and EDA software; Jan 2026 `742.6(b)(10)` case-by-case review for China/Macau chips <21,000 TPP & <6,500 GB/s US26 EAR742.
Frontier AI weights	Proposed `266504X(2)`: >=10^26 FLOPs, >=2T parameters, or Global Top-10 rank at export; Core Data models `266501J` proposed prohibited; APIs outside model export CN-A1 CNQA.	Codified `4E091` (>=10^26 ops closed weights), under May 2025 BIS non-enforcement directive contested by May 2026 GAO Decision `B-337935` US25 USNE GAO26.
Open-source AI	**Footnote 1** excludes already-public open-source AI & papers; **Footnote 2** grants automatic-licensing facilitation ("apply & go") for new open releases & chip tape-outs CN-A1.	**15 CFR 734.7** (published) and **4E091 notes** exclude unrestricted publicly available open-source code and open model weights from licensing EAR734 EAR774.
Online filing portal	`https://ecomp.mofcom.gov.cn` (Technology Trade Management Information Application) CN48 CNPLAT.	`https://snapr.bis.gov` (SNAP-R filing) and `https://snapr.bis.gov/public/stela` (STELA tracker) EAR734 SNAPR.

Consider a hypothetical lab working across China and Singapore that uses covered chips and later licenses private model technology abroad. The chips, parties and uses may engage US rules; the technology transfer may engage Chinese registration. We must check the facts of each activity. One permission does not settle the other system. CN48 EAR734 EAR744

Step-by-Step: How to Register or Apply for Official Approval

Because most writing on export controls stops at ominous warnings, everyday founders, open-source maintainers, university researchers, and equipment buyers are left asking a simple practical question: *"If our release, chip tape-out, or partnership actually falls inside these rules, which government department handles it, what portal do we open, how many working days does approval take, and whom can we call or email?"*

Fig. 13 introduces the documents involved in a filing. The country-specific steps and contacts follow below CN48 CN-A2 CNTIER CNMEAS CNPLAT EAR734 EAR742 SNAPR:



Fig. 13 A researcher presents technical documents for review. The illustration introduces the paperwork, not a complete legal procedure or a promise of approval.

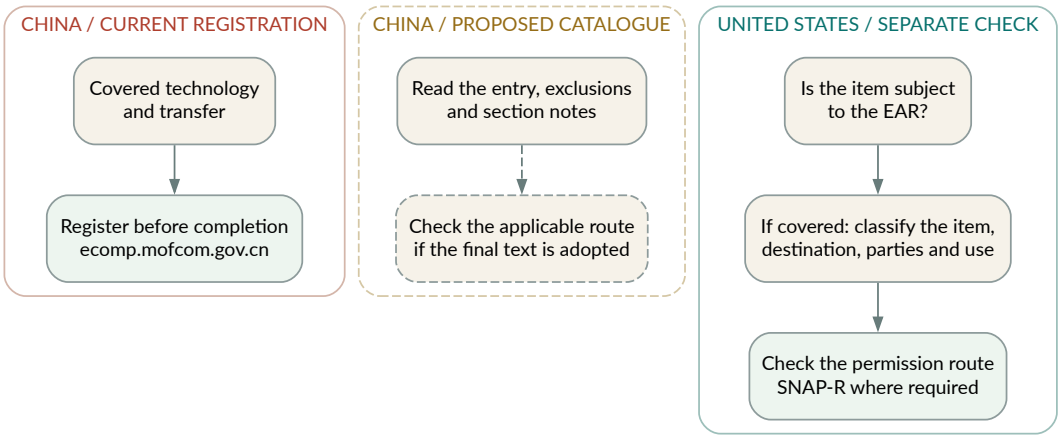


Fig. 14 Current Chinese registration, the proposed Chinese catalogue and US jurisdiction are separate checks. Dashed lines mark the proposed track.

China: Government Departments, Statutory Timelines, and Verified Contacts

In China, technology export administration works as a **central-plus-provincial partnership** under the *Regulations on the Administration of Technology Import and Export* (中华人民共和国技术进出口管理条例, Arts. 30-39) and the *Measures for Prohibited and Restricted Technology Exports* (禁止出口限制出口技术管理办法, MOFCOM & MOST Order No. 2 of 2009, Arts. 4-18): **MOFCOM** and **MOST** set national rules and operate the unified online portal in Beijing, while formal restricted-technology licenses are examined and issued by the **Provincial-level Commerce Department** and **Provincial-level Science and Technology Department** where the exporter is registered CN48 CNTIER CNMEAS.

> **IMPORTANT**
> **Scope of the clocks:** The 30-working-day and 15-working-day periods describe the existing standard restricted-technology procedure in the cited regulations. They are not a guaranteed total of 45 elapsed days, a time limit for temporary registration, or the final procedure for proposed automatic licensing. Negotiation, document preparation and other required reviews can affect the full process. Contact details below are those printed in the cited records; confirm current channels before relying on them.

The Three Official Filing Paths in China

1. **Path A: Mandatory Pre-Export Registration (Active Right Now, since 16 Oct 2026):**
   - **When required:** Before transferring technology ownership abroad, licensing domestic technology for overseas use, or changing a domestic technology owner/actual controller into a foreign entity for any of the 40 points in Attachment 1 (including ongoing unfinished transfers) CN48, section III.
   - **How to file:** Open **MOFCOM's Unified Business System Platform** (`https://ecomp.mofcom.gov.cn`), select the **Technology Trade Management Information Application** (技术贸易管理信息应用), and submit **Attachment 2** (*Technology Export Information Registration Form* with official company seal 公章, owner details, foreign recipient names/country/address/relationship, seven-character code, technical indicators, end use, and USD contract value) CN48 CN-A2 CNPLAT.
2. **Path B: Open-Source, Chip Tape-Out & Cloud Fast-Track ("Automatic Licensing", Proposed in Footnotes 2 & 3):**
   - **When applicable:** Once the revised catalogue takes effect, open-sourcing restricted AI technology, deploying models on a **Chinese firm's self-built, self-operated overseas compute center**, sending chip blueprints abroad for **foundry tape-out** (芯片流片), conducting joint international R&D or standards work, and exporting restricted biomedical targets for out-licensing or humanitarian relief CN-A1, notes 2 and 3.
   - **What is specified:** The draft and official explanation propose automatic licensing or similar facilitation. They do not establish a complete final filing procedure or a guaranteed turnaround for that future route. Confirm the implemented procedure with the competent authority. CN-A1 CNQA
3. **Path C: Standard Two-Stage Restricted Technology Export License (技术出口许可证):**
   - **Stage 1: Before substantive negotiation or signing a contract (up to 30 working days):** File the *Application for Chinese Restricted Export Technology* (中国限制出口技术申请书) with your **Provincial Commerce Department** (省、自治区、直辖市商务主管部门) CNMEAS, Arts. 4-5. Within **5 working days**, Provincial Commerce forwards technical materials to the **Provincial Science & Technology Department** (省、自治区、直辖市科技行政主管部门), which convenes experts to complete technical review within **15 working days**

CNMEAS, Art. 7. Within **30 working days total** from receipt, the authorities decide and, if approved, issue the **Letter of Intent for Technology Export License** (中华人民共和国技术出口许可意向书), valid for **up to 3 years** CNTIER, Arts. 35-36 CNMEAS, Arts. 6, 10. *(Statutory rule: without holding this Letter of Intent, no entity or individual may conduct substantive export negotiations or make legally binding commitments `CNMEAS, Art. 11 <https://www.mofcom.gov.cn/dl/file/20211203230621.pdf>`_.)*

- **Stage 2: After signing the contract (up to 15 working days):** Sign the contract, log contract details in the online system, and submit to Provincial Commerce: **(i)** Letter of Intent, **(ii)** signed contract copy, **(iii)** Technical Documentation Export Checklist (技术资料出口清单), and **(iv)** legal-status proofs of both contracting parties CNTIER, Art. 37 CNMEAS, Arts. 13, 16. Provincial Commerce reviews contract authenticity and decides within **15 working days**, issuing the formal **Technology Export License** (中华人民共和国技术出口许可证) CNTIER, Arts. 37-38 CNMEAS, Art. 14. **Crucial statutory rule:** A restricted technology export contract *takes legal effect only on the date the Technology Export License is issued*, and the exporter presents that license to clear **customs, foreign exchange, banking, and tax** procedures CNTIER, Art. 38 CNMEAS, Arts. 15, 18 CNQA.

Verified China Official Contact Directory

- **Lead Central Departments (Catalogue Policy, Scope & Public Consultation):** Bureau of Industry, Security, Import and Export Control, Ministry of Commerce of the PRC (中华人民共和国商务部安全与管制局), jointly with the Ministry of Science and Technology (科学技术部) CN48.
- **Postal Address:** No. 2 Dongchang'an Avenue, Beijing 100731, China (北京市东长安街2号商务部安全与管制局，邮编：100731) CN48.
- **Official Consultation & Policy Email:** `jinxianmulu@mofcom.gov.cn` *(for public comments through 15 Nov 2026, mark subject* 《中国禁止出口限制出口技术目录》公开征求意见 *and attach contact details and company seal if filing as an organization)* CN48.
- **Official Fax:** `+86-10-6519-7565` (`010-65197565`) CN48.
- **Unified Online Registration & Licensing Portal:** `https://ecomp.mofcom.gov.cn`: *Technology Trade Management Information Application* (商务部业务系统统一平台 - 技术贸易管理信息应用, integrating `jsjckqy.fwmys.mofcom.gov.cn`) CN48 CNMEAS, Art. 16 CNPLAT.
- **Unified Platform Technical Support Phone:** `+86-10-6787-0108` (`86-10-67870108`) CNPLAT.
- **Local Licensing Authorities:** Provincial, Autonomous-Region, or Municipal Commerce Department and Science & Technology Department in the exporter's place of registration (所在地省、自治区、直辖市商务主管部门会同科技行政主管部门) CNMEAS, Art. 4.

United States: Federal Agencies, SNAP-R Workflow, and Verified Contacts

In the United States, dual-use hardware, software, and technology export licensing is administered federally by the **Bureau of Industry and Security (BIS)** within the **U.S. Department of Commerce**, with zero government filing fee for electronic submissions EAR734 EAR742 SNAPR.

Four Steps to Classify or Obtain US Export Approval

1. **Step 1: Establish jurisdiction and publication status:** Check the exact artifact under 734.3, 734.7, 734.8 and the 4E091 notes. Qualifying published material can be outside the EAR; an item that is merely eligible for no license to one destination can still be subject to the EAR. Part 744 requirements generally concern items subject to the EAR, while other sanctions and US-person rules require separate analysis. EAR734 EAR744 EAR774
2. **Step 2: Obtain a Company Identification Number (CIN) on SNAP-R:** Parties exporting controlled hardware (`3A090`), chip equipment, or private controlled technology register online for a CIN on BIS's **Simplified Network Application Process Redesign (SNAP-R)** portal

at `https://snapr.bis.gov` (US-located parties submit Export License applications, while parties abroad may submit Reexport License or Commodity Classification requests) EAR734 SNAPR.

3. **Step 3: File a Commodity Classification (CCATS) or Policy Inquiry if uncertain:** Submit technical specifications as a Commodity Classification Work Item on `snapr.bis.gov` to receive an official ECCN determination (`ACN` tracking number), or contact BIS's Regulatory Policy Division for rule interpretation EAR734 SNAPR.
4. **Step 4: Submit License Application & Track via STELA:** Upload PDF supporting documents, end-user statements, and technical specs (for January 2026 China/Macau chip reviews under `742.6(b)(10)`, attach certifications of `TPP < 21,000`, `DRAM bandwidth < 6,500 GB/s`, US supply adequacy, KYC security, and independent US lab verification) US26 EAR742 SNAPR. Track status on **STELA** (`https://snapr.bis.gov/public/stela`) EAR734.

Verified United States Official Contact Directory

- **Electronic Licensing & Tracking Portals:** SNAP-R Filing Portal: `https://snapr.bis.gov` | STELA Status Tracker: `https://snapr.bis.gov/public/stela` EAR734 SNAPR.
- **SNAP-R Helpdesk & Account Support:** Technical Helpdesk: `snapr@bis.doc.gov` | Account Recovery & Empowered Official Letters: `ECDOEXS@bis.doc.gov` SNAPR.
- **BIS Outreach and Educational Services Division (Washington, D.C.):** U.S. Department of Commerce, 14th and Pennsylvania Avenue, NW., Room H1099D, Washington, DC 20230 | Tel: `+1 (202) 482-4811` | Fax: `+1 (202) 482-2927` EAR734.
- **BIS Western Regional Office (Southern California):** 2302 Martin St., Suite 330, Irvine, CA 92612 | Tel: `+1 (949) 660-0144` | Fax: `+1 (949) 660-9347` EAR734.
- **BIS Northern California Branch (Silicon Valley):** 160 W. Santa Clara Street, Suite 725, San Jose, CA 95113 | Tel: `+1 (408) 998-8806` | Fax: `+1 (408) 998-8677` EAR734.
- **BIS Regulatory Policy Division:** Room 2099B, U.S. Department of Commerce, Washington, DC 20230 | Tel: `+1 (202) 482-2440` | Email: `RPD2@bis.doc.gov` US25 US26 EAR742.
- **BIS Technical Leads (Semiconductors & AI):** Category 3 Semiconductors (`3A090`): Carlos Monroy, `+1 (202) 482-3246`, `Carlos.Monroy@bis.doc.gov` US25 US26 | Category 4 Computers & AI Model Weights (`4E091`): Aaron Amundson, `+1 (202) 482-0707`, `Aaron.Amundson@bis.doc.gov` US25.
- **U.S. Department of the Treasury (Outbound Investment Security Program):** `https://home.treasury.gov/policy-issues/international/outbound-investment-program` | Inquiries: `Outbound@treasury.gov` | Regulations: `OIS.Outbound.Regulations@treasury.gov` | IT Support: `Outbound.ITSupport@treasury.gov` TREAS.
- **U.S. Department of Justice (Data Security Program, National Security Division):** `https://www.justice.gov/nsd/data-security` | Compliance: `NSD.FIRS.datasecurity@usdoj.gov` | Public Affairs: `nsd.public@usdoj.gov` DOJ.

What This Means for the Two Teams

The shared-workbench story now becomes a practical inquiry: what is being transferred, which system covers it, and which route remains available? Clear exclusions, usable procedures and credible continuing cooperation may change incentives. Their effects need evidence; a completed form alone cannot tell us whether cooperation has improved.

The two teams' starting point is in The Fortress Trap: How Both Sides Can Lose.

09. Everyday Situations: A Reader's Field Guide

Can I Still Download DeepSeek, Qwen, or Other Open Models onto My Own Hardware?

Announcement 48 does not establish a blanket ban on downloading existing public models. The Chinese draft excludes already-public AI technology, and US rules contain provisions for qualifying published code and weights. A private fine-tune, restricted data or a later onward transfer needs its own assessment; downloading an original public model is not a universal clearance for every derivative or use. CN-A1, note 1 EAR734 EAR774

Does Hardware Affect Open-Source AI Even When the Software License Is MIT or Apache 2.0?

Profoundly. An MIT license gives you the *legal right* to run a model, whereas physical GPUs, high-bandwidth memory, and cluster interconnects give you the *physical ability* to train or run it. US rules (`3A090`, `742.6(b)(10)`) govern who can buy advanced AI accelerators, while China's proposed `266506X` and `266507X` govern cross-border transfers of AI processor architectures, 3D NAND/DRAM designs, and >10,000-card cluster engineering (with Footnote 2 easing overseas foundry tape-outs) CN-A1 US26 EAR742. Always check both the software files and any custom silicon or cluster stack moving alongside them.

If Overseas Users Chat with Our AI App or Call Our API, Is That a Model Export?

MOFCOM distinguishes remote services from transferring restricted large-model technology. Its answer covers API access, apps, web interaction, SaaS and lightweight local deployment. The explanation is about restricted large models after the measures take effect. It is not a universal exemption for prohibited core-data technology, other controlled items, data rules or every current registration question. CNQA

Our Lab Stays in China, but Foreign Investors Are Acquiring Control. Do We Register?

Check immediately. Section III(2) of Announcement 48 requires pre-export registration whenever the **owner or actual controller of covered domestic technology changes into an overseas**

entity, even if every server and engineer stays put CN48, section III(2). If US investors participate, separately verify US Treasury's Outbound Investment rules (`Outbound@treasury.gov`) TREAS.

What Should We Record in a Two-Page "Release Passport" Before Publishing?

Before any major open-weight release, overseas chip tape-out, or cross-border license, keep a clean internal folder recording: **(1)** exact version and SHA-256 hash; **(2)** catalogue codes checked (`266504X`, `266506X`, `266507X`, `3A090`, `4E091`); **(3)** measured training FLOPs, parameter count, and public benchmark rankings on the release date; **(4)** training-data audit confirming zero Core Data and documenting public/trace Important Data; **(5)** prior publication proof or first-release status; and **(6)** copies of any `ecomp.mofcom.gov.cn` registration or `snapr.bis.gov` ACN filing CN-A1 CN-A2 SNAPR.

Six Real-World Scenarios

> **NOTE**
> **Illustrative scenarios:** The six cases below demonstrate how verified rules operate in practice; always check exact project facts and official primary texts before acting.

1. Fine-Tuning an Existing Public Open-Weight Model
A hypothetical hospital research team obtains public weights and adapts them locally. The original public artifact and a later private clinical adaptation are distinct. Check the applicable publication provisions, data obligations and any onward transfer. A copyright license does not establish clinical safety or grant government permission. CN-A1, note 1 EAR734 EAR774
2. Publishing a Brand-New Top-10 Open-Source Model from China
A Chinese AI lab prepares the very first global open-source upload of a newly trained reasoning model whose live benchmarks place it at #7 in the world (crossing `266504X(2)`). First confirm zero PRC Core Data (`266501J`) and check Important Data trace rules. During the post-16 October 2026 transition window, verify pre-export registration requirements on `ecomp.mofcom.gov.cn`; once the revised catalogue takes effect, use **Footnote 2's automatic-licensing fast lane** ("apply and go") CN48 CN-A1, note 2 CNQA.
3. Licensing Private Frontier Weights or Custom Chip IP to a Foreign Partner
A Beijing team signs a closed commercial contract handing unpublished 2.4-trillion-parameter weights (`266504X`) or a custom 7 nm NPU instruction-set core (`266506X`) to an overseas enterprise. Because this is a private commercial transfer rather than open-source publication, file **Attachment 2 pre-export registration** on `ecomp.mofcom.gov.cn` today, and prepare for the **two-stage Provincial Commerce + Provincial S&T license** (30 working days for Letter of Intent + 15 working days for formal Technology Export License) once finalized CN48 CNTIER CNMEAS. *(Note: sending a chip layout abroad solely for foundry wafer fabrication qualifies for Footnote 2 facilitation.)*
4. Offering Global Web Chat or Metered API Access
A Shanghai startup serves overseas customers through a web interface and HTTPS API powered by its flagship frontier model hosted in China. MOFCOM's official Q&A confirms that remote API and app services without transferring underlying weights or source code sit **outside large-model technology export controls** (provided the model was not trained on prohibited Core Data `266501J`) CN-A1 CNQA.
5. Foreign Venture Restructuring or Cross-Border Acquisition
A domestic startup holding unreleased `266505X` 60-agent orchestration software or an undisclosed `262705X` gene-therapy target discovery engine restructures so an offshore holding company

becomes its actual controller. Even if zero files leave Shenzhen servers, Section III(2) of Announcement 48 requires **pre-export registration before the control shift completes**, and if US capital is involved, US Treasury's Outbound Investment rules must be checked in parallel CN48, section III(2) TREAS.

6. Shipping a Turnkey Solar Wafer or Perovskite Line Overseas

A Chinese equipment maker contracts to ship monocrystalline crystal-pulling furnaces and PECVD coaters to equip a new overseas solar plant, bundled with process recipes and engineer commissioning. Under **Footnote 4**, the physical machines themselves fall under export control if they embody listed `263004X-263802X` indicators (and `<=70 micrometer` wafer slicing `263004J` or `>=30%` tandem recipes `263801J` face proposed prohibition). Complete `ecomp.mofcom.gov.cn` pre-export registration before completion CN48 CN-A1, note 4 CNQA.

10. Primary Sources, Method, and Plain Glossary

How Every Fact in This Book Was Verified

We checked the claims against 26 legal, technical and research sources as of 2 November 2026. Added research informs the discussion of cooperation and risk; it does not turn our assumptions into measured results. Contact information is what the cited records print and may change. Source verification is distinct from approving a transaction.

Original WPS bytes and SHA-256 hashes are preserved untouched alongside converted PDF reading copies. Where GAO's server blocked automated byte archiving (`HTTP 403`), we note that limitation transparently. Throughout the book, primary legal facts carry bracketed citations, while our game-theoretic matrix and future scenarios are presented openly as analytical models rather than legal advice for a specific transaction.

Untouched Official Document Copies

- Announcement No. 48 Official Page (PDF reading copy)
- Attachment 1: 40 Proposed Technologies (Original WPS) | Converted PDF
- Attachment 2: Pre-Export Registration Form (Original WPS) | Converted PDF
- MOFCOM & MOST Order No. 2 of 2009: Licensing Measures (Official PDF)
- US BIS Jan 2026 Semiconductor Rule 91 FR 1684 (Official PDF)
- US BIS Jan 2025 AI Diffusion Rule 90 FR 4544 (Official PDF)
- US BIS Official SNAP-R Filing Guide & FAQ (Official PDF)
- Machine-Readable Provenance & SHA-256 Manifest (sources.json)

Plain-Language Glossary for Normal Humans

Model weights (parameters)
: The billions or trillions of tuned numerical values learned during training that shape a model's outputs. Weights alone do not include every part of a reproducible or safe system.

Checkpoint
: A saved snapshot of weights during or at the end of training, used to resume training, distil smaller models, or fine-tune for specialist tasks.

FLOPs (floating-point operations)
: Individual math steps performed during training. 10^{26} **FLOPs** (a 1 followed by 26 zeros) is the frontier compute threshold in both China's proposed `266504X` and America's `4E091`.

TPP & HBM (the hardware bottleneck)
: **Total Processing Performance (TPP)** measures an AI chip's raw compute throughput,

while **High-Bandwidth Memory (HBM / DRAM bandwidth)** measures how fast data feeds into the chip (capped at `TPP < 21,000` and `6,500 GB/s` in the January 2026 US China-review rule `742.6(b)(10)`).

Chip tape-out (芯片流片)
Sending a finished processor blueprint from a chip-design team to a semiconductor foundry to manufacture physical silicon wafers, explicitly offered automatic-licensing facilitation in Announcement 48's Footnote 2.

API / SaaS (the bakery window)
Letting remote users send prompts to your cloud server and receive answers back, without handing over underlying model weights, source code, or chip designs.

Nash equilibrium
A set of choices in which no player can improve its own payoff by changing only its own choice. An equilibrium need not be a stalemate, safe, fair or socially best. Mutual walls are an equilibrium only under the assumptions of our example. GAME

Software license vs. export license
An *MIT* or *Apache 2.0* license is a private copyright permission from creators allowing anyone to reuse code or weights. A *Technology Export License* (`技术出口许可证` or BIS license) is a sovereign government permit allowing controlled technology or hardware to cross a national border.

Automatic licensing (自动许可)
The streamlined "apply-and-go" (即报即出) government permit channel proposed in Announcement 48's footnotes for open-sourcing restricted AI, self-owned overseas clouds, overseas chip tape-out, and humanitarian biomedicine.

ECCN, CCATS & SNAP-R
In the US system, an **ECCN** (`3A090` for AI chips, `4E091` for weights) identifies a controlled item; **SNAP-R** (`snapr.bis.gov`) is the free electronic portal where exporters file license applications and **CCATS** classification requests.

Complete Register of 26 Verified Primary Sources

Table 4 Primary legal, regulatory, and technical sources verified 2 November 2026

Code	Official Source & Scope Verified	Date	Primary URL & SHA-256 Prefix
`CN48`	MOFCOM & MOST, Announcement No. 48 of 2026 (Sections I, II & III)	2026-10-16	mofcom.gov.cn (`a5261a74...`)
`CN-A1`	Announcement 48, Attachment 1: 40 Proposed Technologies & 4 Footnotes (WPS)	2026-10-16	Original WPS (`213158bb...`)
`CN-A2`	Announcement 48, Attachment 2: Pre-Export Registration Form (WPS)	2026-10-16	Original WPS (`d65818c2...`)
`CNQA`	MOFCOM Spokesperson Official Q&A on Announcement No. 48	2026-10-16	mofcom.gov.cn (`37fc685a...`)
`CNTIER`	PRC State Council, Regulations on Administration of Tech Import & Export (Arts. 30-39)	2020-11-29	xzfg.moj.gov.cn (`39f4b9bb...`)
`CNMEAS`	MOFCOM & MOST Order No. 2 of 2009, Measures for Prohibited/Restricted Tech Exports	2009-04-20	mofcom.gov.cn PDF (`d14e5230...`)

Code	Official Source & Scope Verified	Date	Primary URL & SHA-256 Prefix
`CNPLAT`	MOFCOM Unified Business System Platform Notice & Support (`+86-10-6787-0108`)	2026-11-02	mofcom.gov.cn (`5bc9740f...`)
`US26`	US BIS Final Rule on Semiconductor License Review (`91 FR 1684`; `742.6(b)(10)`)	2026-01-15	govinfo.gov (`d3b27863...`)
`US25`	US BIS Framework for Artificial Intelligence Diffusion (`90 FR 4544`; `4E091`)	2025-01-15	govinfo.gov (`6e202458...`)
`USNE`	US BIS Press Release Directing Non-Enforcement of AI Diffusion Rule	2025-05-13	bis.gov (`14a05e6e...`)
`GAO26`	US GAO Decision `B-337935` on AI Diffusion Rule & Congressional Review Act	2026-05-12	gao.gov *(HTTP 403; web-verified)*
`EAR734`	15 CFR Part 734: Scope of the EAR, Published Open Source (`734.7`) & Contacts	2026-11-01	bis.gov/ear/734 (`eeb76281...`)
`EAR742`	15 CFR Part 742: Control Policy, CCL Based Controls (`742.6(b)(10)`)	2026-11-01	bis.gov/ear/742 (`3eccf859...`)
`EAR744`	15 CFR Part 744: End-User and End-Use Based Controls (`744.11`, `744.23`)	2026-11-01	bis.gov/ear/744 (`ab19c823...`)
`EAR774`	15 CFR Part 774: Commerce Control List (ECCNs `3A090` and `4E091`)	2026-11-01	bis.gov/ear/774 (`01d3beba...`)
`SNAPR`	US BIS SNAP-R Electronic Filing FAQ & Contact Guide (`snapr.bis.gov`)	2026-11-02	snapr.bis.gov PDF (`e0c3bdee...`)
`TREAS`	US Treasury Outbound Investment Security Program & Oct 2026 Proposed Rule	2026-10-22	home.treasury.gov (`75f52eee...`)
`DOJ`	US Department of Justice National Security Division: Data Security Program	2026-11-02	justice.gov/nsd (`1bf01adb...`)
`OSI`	Open Source Initiative: Open Source AI Definition 1.0	2024-10-28	opensource.org (`a5442179...`)
`R1`	DeepSeek-R1 Official Repository & MIT License Record	2025-01-20	github.com/deepseek-ai (`875b282c...`)
`QWEN`	Qwen3 Official Repository & Apache 2.0 License Record	2025-04-29	github.com/QwenLM (`2b7cdab9...`)
`GD`	GuideDog Documentation & Book Publishing Engine (`v1.0.0-beta.10`)	2026-11-02	github.com/insanai/guidedog (`6fe7704f...`)

Research and Historical References Added in This Review

- `NTIA`: Dual-Use Foundation Models with Widely Available Model Weights (July 2024). A government assessment of benefits, risks and monitoring recommendations for then-current open models, not evidence about every future model.
- `OPENRISK`: On the Societal Impact of Open Foundation Models. A research position paper on marginal risks and benefits. We distinguish its argument from a measured causal effect of export controls.
- `GAME`: A Course in Networks and Markets, Chapter 1. Definitions of games, best responses and equilibrium. Our particular scores and repeated-game calculation are the book's own example.
- `JEFFERSON`: Thomas Jefferson to Isaac McPherson, 13 August 1813. The original source of the candle quotation. The analogy does not itself establish an economic effect.

Credits and Reuse

Written by **Vikrant Rathore**, with assistance from **Jennifer Yuan, Vicky Ni, Ronak Rathore and Kanak Rathore**.
Original text, editorial work and new illustrations are licensed under **CC BY-SA 4.0**. Attribution and sharing adaptations under the same license are required. Official documents, quoted material, fonts, software and other third-party materials retain their own copyrights and licenses. See the repository's LICENSE and NOTICE files for the terms and exclusions.
All seven language editions are editorial translations. They have not received independent professional legal or native-language review. Original legal texts prevail. Examples and game scores must not be treated as official classifications or empirical forecasts.

Optional: Checking the Game Example

The opening story can be understood without these calculations. Here we make the extra assumptions explicit for readers who want to check them.

Why the Ending Is Stable

The ranking 20 > 17 > 5 > 3 gives each country a reason to restrict whatever the other chooses. With two walls, either country sharing alone would move its score from 5 to 3. Mutual sharing is Pareto-efficient in this small game: no other listed outcome improves one country's score without lowering the other's. That does not mean either country receives its largest individual score. GAME

When Future Cooperation Matters

For the next calculation, we additionally treat the invented scores as comparable payoffs in each round. The numerical threshold depends on that scale, not just on the order of preferences.
Countries also interact repeatedly. In this simplified game, suppose cooperation continues while both share, but a deviation earns 20 once and is followed by 5 forever. Let δ represent how much a country values the next round relative to this one. Cooperation is at least as attractive as that deviation when:

$$\frac{17}{1-\delta} \geq 20 + \frac{5\delta}{1-\delta} \quad \Longleftrightarrow \quad \delta \geq \frac{1}{5}$$

This is a calculation inside the assumed game, not a diplomatic forecast. It assumes an indefinite relationship, observable choices and a credible response to deviation. Errors, hidden actions or a known final round can undermine that strategy. The useful policy question is how to make continued cooperation credible, rather than treating mistrust as inevitable. GAME

Optional: Independent Review Is Not a Safety Guarantee

Open access can make independent testing and repair easier. A simple probability example explains one possible benefit. If a flaw exists with probability p_0 and each of n teams independently detects it with probability q, then:

$$P(\text{flaw remains undetected}) = p_0(1-q)^n$$

For a flaw known to exist and an assumed detection chance of one half, one team misses it with probability 50%, four independent teams with 6.25%, and ten with about 0.098%. These are conditional calculations. Real teams can share blind spots, lack testing resources or find a flaw without fixing it. An undetected-flaw probability is not a measure of total cyber, biological or societal risk. Open weights can also help harmful actors modify a system or remove safeguards. NTIA's 2024 report considered both benefits and risks, recommending monitoring rather than restricting then-current open weights. Research on open foundation models stresses the need to compare the

extra risk with existing alternatives. Neither source establishes a universal safety multiplier for openness. NTIA OPENRISK