

Application of U.S. Export Controls to AI Biosecurity Evaluations and Mitigations

By Doni Bloomfield, Joe Khawam, and Tim Schnabel

October 14, 2025

Executive Summary: To ensure that frontier AI models cannot be exploited to create or enhance biological threats, rigorous evaluations of such models are essential, as are effective mitigations of any dangerous capabilities revealed by the evaluations. Across U.S. policy circles and industry practice, pre-deployment biosecurity evaluations—paired with commensurate mitigations—are increasingly treated as a baseline prerequisite for responsible model release.¹ The Trump administration’s AI Action Plan also recognizes the importance of biosecurity evaluations, highlighting the need to invest in measures to evaluate and mitigate risks posed by advanced AI models.²

However, these evaluations and mitigations involve protocols for creating or modifying pathogens that may implicate export controls. Data used in or generated by such protocols can constitute “technical data” controlled by the International Traffic in Arms Regulations (ITAR) or “technology” controlled by the Export Administration Regulations (EAR). Compounding the challenge, only a small number of experts possess the requisite biosecurity expertise to perform such evaluations and design effective mitigations, and some of these experts are foreign nationals. In a competitive landscape, the window for fully testing a new model, incorporating necessary safeguards, and deploying it is only a few weeks, starting when the model is ready to be evaluated. By contrast, obtaining a case-by-case license from the Directorate of Defense Trade Controls (DDTC) or Bureau of Industry and Security (BIS) can take a month or longer.³ The timing mismatch—exacerbated by the difficulty of distinguishing controlled from uncontrolled data—risks discouraging model developers and evaluators from seeking approvals, or driving them to curtail testing and mitigation so as not to slow down the release of new models.

This paper maps those export-control hazards, surveys existing regulatory tools, and considers additional actions for reconciling compliance with urgent biosecurity needs. It concludes that, when only the EAR is implicated, most obstacles (other than those related to ricin and saxitoxin) can be mitigated by limiting 1E/2D/2E disclosures to U.S. persons and employees and evaluators who are nationals of Australia Group countries. When ITAR-controlled technical data is involved, developers and evaluators must either (1) restrict access to U.S. persons in the United States or (2) if foreign person access is necessary, pursue a Technical Assistance Agreement or other appropriate authorization with DDTC.

¹ See Frontier Model Forum, *Issue Brief: Preliminary Taxonomy of AI-Bio Safety Evaluations* (Dec. 20, 2024), <https://www.frontiermodelforum.org/updates/issue-brief-preliminary-taxonomy-of-ai-bio-safety-evaluations/>.

² See White House, *Winning the Race: America’s AI Action Plan* 10, 22, 23 (2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

³ See Directorate of Def. Trade Controls Public Portal, http://www.pmdtcc.state.gov/ddtc_public/ddtc_public?id=ddtc_public_portal_homepage (indicating monthly averages of between 25 and 56 days of process time for DCS licenses between August 2024 and August 2025); Bureau of Industry and Security, *Fiscal Year 2023 Annual Report* 7 (2024), <https://www.bis.doc.gov/index.php/documents/about-bis/newsroom/3569-bis-fiscal-year-2023-annual-report/file> (stating that “BIS’s average license application processing time in FY 2023 was 38 days”).

While existing regulations and tools may be sufficient in many cases to meet industry needs, BIS and DDTC should consider two further actions to reduce the uncertainty and burden of export control compliance with respect to AI biosecurity evaluation and mitigation activities.

First, BIS and DDTC should publicly identify the licensing vehicles that are presumptively appropriate when authorization is needed to perform AI biosecurity activities. BIS should also clarify that U.S. person controls do not require entities to seek a license to perform bona-fide AI biosecurity evaluations. The agencies could provide this guidance in the form of concise and direct FAQs that focus on current regulatory uncertainties rather than speculative scenarios that could constrain future innovation or adaptation to rapidly evolving technologies. To address the timing challenges identified above, the agencies could also establish expedited review procedures for biosecurity license requests. BIS and DDTC should work closely with developers and evaluators to ensure that the guidance and procedures will have the intended effect of encouraging evaluation and mitigation activity.

Second, BIS and DDTC should also consider rulemaking to adopt a narrow EAR license exception with a mirror ITAR exemption for bona-fide AI biosecurity evaluations and mitigations. Such a carve-out—with appropriate conditions—would allow developers and evaluators to collaborate more easily with foreign specialists while preserving non-proliferation safeguards. Together, targeted guidance and a limited carve-out would synchronize regulatory timelines with industry practice, supporting the voluntary evaluation ecosystem on which safe deployment of frontier AI models increasingly depends.

This paper analyzes export-control implications in the context of *biosecurity* evaluations of frontier AI systems as an especially urgent need. Developers often conduct additional chemical, radiological, and nuclear evaluations during the same campaigns, which may implicate distinct provisions of the EAR and ITAR, as well as additional regimes, such as the Nuclear Regulatory Commission's (NRC) controls. This paper does not analyze those domains of evaluation or related mitigations, though they remain critical. BIS and DDTC—working in coordination with NRC and other relevant agencies—may therefore wish to address these additional areas in parallel with biosecurity evaluations.

I. Why are AI biosecurity evaluations and mitigations important?

Frontier LLMs, specialized biological models, and generative-AI agents may lower barriers to creating or modifying novel, extinct, or enhanced pandemic-potential pathogens. Automated evaluations, red-teaming exercises, expert challenge rounds, and corresponding mitigations are currently the primary practical ways to detect and reduce bio-leakage risks before public

deployment—and must recur as models evolve.⁴ Appendix 1 provides a brief overview of a typical workflow for these processes, and Appendix 2 provides some examples of specific evaluations from one recent model’s system card.

A limited number of specialists possess the biological threat-assessment expertise to design and implement these evaluations and mitigations, and some of those specialists are foreign persons. Due to high demand for such expertise, U.S. developers and evaluators often must engage foreign experts as contractors, hire foreign person employees, or hire U.S. persons living abroad. If U.S. firms disclose controlled information to such persons, the disclosures would be subject to export-control obligations.⁵

II. What export control issues arise from AI biosecurity evaluations and mitigations?

Export-control hazards may arise in various ways during, or in preparation for, AI biosecurity evaluations and mitigations:

- **Deemed exports to foreign persons** inside the United States occur when foreign contractors or employees view or discuss controlled information.
- **Deemed reexports to foreign persons** occur when foreign contractors or employees located outside of the United States view or discuss controlled information.
- **Exports of controlled information** occur when controlled information is released to persons, including U.S. persons, outside the United States.
- **Aggregation of dangerous know-how into test corpora** for adversarial prompting can create a controlled dataset subject to the ITAR or EAR. Simply compiling publicly available information does *not* by itself create controlled data; however, if the aggregation adds controlled know-how or synthesizes the material in a manner that generates information required for the development, production, or use of ITAR-controlled defense articles or EAR-controlled items, the relevant portions of the dataset may be subject to the ITAR or EAR.
- **Provision of defense services** (under the ITAR) may occur when U.S. person employees or evaluators furnish technical data to foreign personnel or furnish assistance to foreign persons in the use of technical data.
- **Uploading controlled data to servers abroad**—or to U.S.-based servers accessed by foreign persons—may constitute an intangible export in certain circumstances.

These hazards may be encountered at several points during the type of evaluation and mitigation process outlined in Appendix 1, including in the context of sharing data sets, prompts, outputs and evaluation results.

⁴ See generally Forecasting Research Institute, *Forecasting LLM-Enabled Biorisk and the Efficacy of Safeguards* (July 1, 2025), <https://forecastingresearch.org/s/ai-enabled-biorisk.pdf>.

⁵ A similar obstacle arises if non-U.S. frontier developers—such as Mistral—were to seek to hire U.S. firms or nationals to red-team or certify their models. The transfer of pathogen-related testing prompts or testing results back to the foreign client may constitute an export or defense service under the ITAR and/or EAR, triggering DDTC and/or BIS licensing requirements. These compliance frictions discourage foreign customers from retaining U.S. evaluators, eroding U.S. competitiveness in this emerging market and reducing the U.S. government’s leverage over global standards.

When licenses are required, BIS and DDTC case-by-case licensing can take a month or more—which can be substantially longer than the few weeks typically available for biosecurity evaluations and mitigations. Strong industry competition and pressure for rapid model releases create challenges for such activities.

Moreover, some industry participants remain uncertain about—or even unaware of—possible export-control obligations relating to biosecurity evaluations and mitigations. Ambiguities may arise around (1) whether certain data is controlled under the ITAR, the EAR, or neither, and (2) how aggregating public-domain information with model-generated reasoning might convert otherwise uncontrolled data into controlled data.

Thus, labs, their contractors, and independent companies or non-profits may limit or omit critical biosecurity evaluations and mitigations, or they may undertake such activities despite potential export control violations or ambiguities. Neither of these outcomes is desirable from a national security and regulatory perspective.

III. How does the EAR generally limit biosecurity AI evaluations and mitigations?

Under the EAR, Commerce Control List (CCL) controls apply to certain algorithms, data, or code that are required for the design, synthesis or enhancement of a controlled pathogen or toxin, or related genetic elements; information required for the disposal of controlled pathogens, toxins, and related genetic elements; and information required for the development, production, or use of certain laboratory equipment or software.

- **ECCN 1C351, 1C353, 1C354:** Controls specific human, animal, and plant pathogens, toxins, and some related genetic elements.
- **ECCN 1E001:** Controls technology required for the development or production of items on the CCL, including 1C351, 1C353, and 1C354.⁶
- **ECCN 1E351:** Controls technology required “for the disposal” of “microbiological materials” described in ECCN 1C351, 1C353, and 1C354.
- **ECCN 2B352:** Controls various forms of equipment capable of handling biological materials, including many types of biosafety items as well as nucleic acid assemblers and synthesizers meeting particular specifications.
- **ECCN 2D352:** Controls software designed for nucleic acid assemblers and synthesizers controlled by 2B352.j that is capable of designing and building functional genetic elements from digital sequence data.
- **ECCNs 2E001, 2E002, 2E301:** Control technology required for the development of items on the CCL, including 2B352 and 2D352, and for the production or use of items including 2B352.

All 1E/2D/2E exports to “CB:2” destinations—states outside the Australia Group (e.g., China, Israel, Russia)—require a BIS license. Technology required for the production of ricin and saxitoxin (e.g., under 1C351.d.15 and d.16), as well as their genetic elements and methods of their disposal, are controlled for release world-wide. General Prohibition 5, along with Part 744, blocks

⁶ Technology is controlled if it is “required” for the production, development, or use of an item, i.e., it is “peculiarly responsible for achieving or exceeding the controlled ... characteristics.” 15 C.F.R. pt. 774, supp. 2; 15 C.F.R. pt. 772.

any unlicensed export, reexport, or transfer to certain restricted parties or end uses, while General Prohibition 6 extends the same bar to destinations under the country embargoes in Part 746.⁷

Under General Prohibition One, it is generally unlawful to export, reexport, or transfer any controlled item to controlled locations without authorization.⁸ “Exports” include releasing controlled technology or source code to a foreign person in the U.S. (deemed export) or a person abroad.⁹

At the same time, information and software that are “published” or arise during or from “fundamental research” are not subject to the EAR.¹⁰ Three license exceptions may also potentially be relevant in some instances:

- **GOV:** Authorizes exports for U.S. Government personnel and agencies and certain exports by, for, or at the direction of DoD or DoE.¹¹
- **TSU:** Authorizes exports and reexports of specific categories of technology and software, including the release of most technology and source code in the United States by U.S. universities to their bona-fide and full-time regular employees other than nationals of arms-embargoed nations (e.g., China, Russia).¹²
- **STA:** Authorizes exports, reexports, and transfers of software source code and technology to specific destinations, but may not be used for items controlled by ECCNs 1C351.a, .b, .c, d.14, d.15, .e, 1C353, 1C354, 1E001 (as it relates to those earlier ECCNs), or 1E351.¹³

The EAR also prohibits U.S. persons from performing, without a license, any service they know “may assist or benefit” the design, development, or production of biological weapons anywhere in the world.¹⁴ These controls should not be seen as requiring a license before performing safety evaluations that mitigate rather than support the creation of biological weapons.¹⁵

IV. How would BIS currently authorize biosecurity evaluation and mitigation activities?

When authorization is required for an export, BIS can issue a single-transaction license. However, individual licenses typically take a month or more to process—potentially far longer than the few-

⁷ 15 C.F.R. §§ 736.2(b)(5), (b)(6), pts 744 & 746.

⁸ 15 C.F.R. § 736.2(b)(1).

⁹ 15 C.F.R. § 734.13(a)(1)-(2), (b).

¹⁰ 15 C.F.R. §§ 734.3, 734.7, 734.8.

¹¹ 15 C.F.R. § 740.11.

¹² 15 C.F.R. § 740.13.

¹³ 15 C.F.R. § 740.20.

¹⁴ 15 C.F.R. § 744.6.

¹⁵ This conclusion holds for two reasons. First, the creation or design of a pathogen or delivery system for safety testing purposes is very likely not the design or production of a biological weapon. Although the EAR does not define “biological weapons,” the term is defined elsewhere in the U.S. Code, and the Biological Weapons Convention of which the U.S. is a party, to mean biological agents or toxins (or parallel delivery systems) developed “other than [for] prophylactic, protective, bona fide research, or other peaceful purposes.” 18 U.S.C. § 175; *see also* Convention on the Prohibition of the Development, Production, and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on Their Destruction art. I, Apr. 10, 1972, 26 U.S.T. 583, 1015 U.N.T.S. 163. Second, although “may assist or benefit” is capacious language, interpreting it to include activities that are meant in good faith to reduce the chances that biological weapons are created would likely render the regulation unconstitutionally vague.

week window in which frontier-model evaluations and mitigations generally must be scoped, executed, and closed.

Under the EAR, BIS can also issue multi-year authorizations—documented as technology licenses or authorization letters. To obtain such an authorization, applicants submit a Technology Letter of Explanation meeting the requirements in Supplement No. 2 to Part 748, paragraph (o) and § 748.8(o). BIS then issues an authorization letter permitting repeated releases of specified “technology” or “software” to named foreign parties. These licenses are valid for four years by default (§ 750.7(g)), though BIS may approve longer—or shorter—terms. Additionally, BIS exercises discretionary authority to issue Supply-Chain Authorization Letters (SCALs), which provide company-specific relief for ongoing or urgent supply-chain transactions. These authorizations may remain valid for several years, subject to conditions specified in the letter.

V. What is the practical effect of the EAR on AI biosecurity evaluations and mitigations?

A disclosure of 1E/2D/2E technology or software may constitute an export and, absent a BIS license or applicable exception, may not be shared with foreign persons from CB:2 destinations, whether those persons are located inside or outside the United States. Releasing controlled technology to restricted parties or nationals from embargoed destinations is also prohibited.

Private information controlled under the EAR may not qualify for the publication or fundamental-research exclusions.¹⁶ In addition, license exceptions GOV, TSU, and STA may apply in narrow circumstances, but none covers the full range of AI biosecurity evaluation and mitigation activities.

However, many foreign biosecurity evaluators and foreign employees are from Australia Group nations, such as Canada, the UK, and the EU member states. No license is required so long as such evaluators and developers disclose 1E/2D/2E technology or software only to foreign persons from these countries and technology related to ricin and saxitoxin is not included.¹⁷ If that is not possible, developers and evaluators engaging in AI biosecurity evaluations and mitigations currently must assess the applicability of the limited license exceptions, omit the controlled technology from their evaluation and mitigation activities, or seek a BIS license.

BIS has not provided guidance on whether model weights themselves can be implicated by export controls based on the information that the weights can produce when queried, or whether any such controls are affected by the extent to which the model is designed not to divulge such

¹⁶ Outputs from AI models trained on public datasets may of course include “published” information under the EAR or “public domain” information under the ITAR. However, AI model outputs may not be “published” or “public domain” simply because the AI model is trained on public datasets. AI models can create novel combinations and insights based on public information that could potentially constitute new controlled information. In addition, models often train on both public and proprietary datasets, which—even if developers seek to avoid any specific datasets known to contain controlled data—may further enhance the ability of models to generate technical outputs that exceed any individual training source. Finally, the method through which information enters the public domain is relevant to determining whether it is “public domain” information excluded under the ITAR, *see* 22 C.F.R. § 120.34(a), and the EAR also continues to control certain limited categories of published information. 15 C.F.R. § 734.7(b), (c).

¹⁷ Whether there are additional technology controls to all destinations under the EAR, e.g. for “CW” reasons, is beyond the scope of this paper.

information.¹⁸ It is thus an open question whether, if a proprietary AI model can output controlled information under 1E/2D/2E, the weights may be controlled to CB:2 destinations and persons, like other forms of controlled 1E/2D/2E software or technology (or to all destinations and persons in the case of ricin or saxitoxin).

VI. How does the ITAR generally apply to biosecurity AI evaluations and mitigations?

Defense articles controlled under the U.S. Munitions List (USML) include biological agents and toxins and directly related technical data.

- **Category XIV(b):** Controls specific genetically modified pathogens that are non-naturally altered to increase pathogenicity, persistence, or to defeat detection/immunity.¹⁹
- **Category XIV(g):** Controls polynucleotides and recombinant materials (e.g. DNA sequences, plasmids, viruses, expression vectors) specific to Cat. XIV agents.
- **Category XIV(m):** Controls technical data (defined by 22 C.F.R. § 120.33 to include the information required for the design, development, production, or testing of defense articles and software directly related to defense articles) directly related to Cat. XIV articles.

The ITAR prohibits the export of defense articles or technical data, or the furnishing of defense services, without DDTC authorization.²⁰ An “export” includes the release or transfer of technical data to a foreign person in the U.S. (a “deemed export”) or the transmission of technical data out of the United States.²¹ A subsequent release of that same technical data abroad to a foreign person who is a citizen of a country other than where the release takes place is a “reexport,” while a “retransfer” occurs if the release is to a foreign person who is a citizen of the country where the release occurs.²² Separately, furnishing to foreign persons any technical data, or furnishing assistance to foreign persons in the design, development, production, or testing of defense articles (including technical data), whether in the U.S. or abroad, constitutes a “defense service.”²³

The ITAR establishes a policy of denying licenses for exports of defense articles and defense services to or from certain prohibited countries, including China.²⁴

Two general exclusions from the ITAR are potentially relevant here:

¹⁸ Cf. 15 C.F.R. § 734.18(a)(5) (transmitting unclassified technology or software, if sufficiently encrypted, is generally not an export).

¹⁹ Category XIV(f) of the USML controls equipment “specially designed” for (1) the dissemination and dispersion of articles controlled in Cat. XIV(b), or (2) testing the articles controlled in Cat. XIV(b) and developed under a DoD contract or other funding authorization. The requirement of a DoD contract or other funding authorization means that the latter control is unlikely to be applicable to private sector biosecurity testing of AI models.

²⁰ 22 C.F.R. § 127.1.

²¹ 22 C.F.R. § 120.50.

²² 22 C.F.R. §§ 120.51, .52.

²³ 22 C.F.R. § 120.32. DDTC has issued a proposed rule (89 FR 60980) in July 2024 that would remove furnishing technical data to a foreign person from the definition of “defense service.” Furnishing such data to a foreign person is already a controlled event, consistent with the definition of an export, reexport, and retransfer, as described above. The proposed rule maintains within the definition of “defense service” the furnishing of assistance to foreign persons in the development (including design), production, or testing of a defense article.

²⁴ 22 C.F.R. § 126.1.

- **Public-Domain Exclusion:** Excludes “public domain” information, which is defined to include information that is published and generally accessible or available to the public through specified activities, including conference releases, patents, or “fundamental research” in science and engineering at universities.²⁵
- **General Science-Instruction Exclusion:** Technical data is defined to exclude information concerning general scientific, mathematical, or engineering principles commonly taught in schools.²⁶

Additionally, four exemptions may occasionally be relevant but generally will not apply in this context:

- **Intra-Company DN/TCN:** Authorizes foreign entities who are authorized end-users or consignees under an existing ITAR authorization to transfer unclassified defense articles (including technical data) to their dual/third-country national employees without further DDTC approval, provided the entity maintains a technology security/clearance plan, screens employees for substantive contacts with § 126.1 countries, executes non-disclosure agreements, and retains screening records for five years.²⁷
- **Canada:** Permits license-free exports of unclassified defense articles to Canadian-registered persons or government authorities for Canadian end-use, except for items listed in Supplement No. 1 to Part 126 (note: Cat. XIV(b) and related technical data are excluded for Canada), provided recipients have no ties to § 126.1 destinations and transfers remain within Canada with five-year record retention.²⁸
- **AUKUS:** Allows license-free defense trade between pre-approved “authorized users” for most unclassified USML items, with § 126.7 enabling trilateral trade among the U.S., UK, and Australia, and §§ 126.16-126.17 enabling bilateral trade between U.S.-Australia and U.S.-UK respectively.²⁹ The exemptions exclude items listed in Supplement No. 1 to Part 126 (for bilateral treaties) and Supplement No. 2 to Part 126 (for AUKUS); Cat. XIV(b) and related technical data are excluded under Supplements No. 1 and 2 for the UK and Australia. Authorized users must comply with respective national security standards and maintain records.
- **U.S.-Government Use:** Covers exports and defense services by or for U.S. Government agencies in furtherance of official programs or contracts, including private contractor activities when performed under government contract authority that specifically directs such exports.³⁰

VII. How would DDTC currently authorize biosecurity evaluation and mitigation activities?

If no exclusions or exemptions apply, DDTC can issue a single-transaction license. As with BIS, individual licenses typically take a month or more to process.

²⁵ 22 C.F.R. § 120.34.

²⁶ 22 C.F.R. § 120.33(b).

²⁷ 22 C.F.R. § 126.18.

²⁸ 22 C.F.R. § 126.5.

²⁹ 22 C.F.R. §§ 126.7, 126.16-.17.

³⁰ 22 C.F.R. §§ 125.4(b)(3), 126.4.

For collaborations that involve repeated transfers of ITAR-controlled technical data or the ongoing provision of defense services, DDTC generally approves a Technical Assistance Agreement (TAA) rather than multiple one-time licenses. A TAA—governed by Part 124 of the ITAR and defined in § 120.57(e)—allows a registered U.S. exporter to furnish controlled data or services, including intangible transfers to specifically named foreign licensees under a single authorization that can remain valid for up to ten years.

A TAA is an executed agreement between the registered U.S. exporter of the technical data and every foreign licensee. When seeking approval of the TAA from DDTC, the ITAR requires a transmittal letter with prescribed information and a statement of work that describes the controlled data and activities in reasonable detail. Ongoing duties include amending the TAA when adding new parties, nationalities, or USML content; filing annual sales reports; and retaining all export and service records for five years. Because DDTC’s Guidelines for Preparing Agreements elaborate on these obligations at length, many exporters rely on specialized counsel or compliance professionals to draft, file, and administer a TAA.

A TAA is generally used when controlled technical data is transferred to foreign person third-party contractors on a continuous basis. If, however, the recipient meets the definition of a “regular employee” in 22 C.F.R. § 120.64—which covers permanent hires as well as long-term full-time contractors under the company’s direction—DDTC generally issues a stand-alone DSP-5 license instead of a TAA. In either case, DDTC has authority under the ITAR to authorize exports of technical data from U.S. persons to foreign person employees or contractors in a prospective manner on a continuous basis.

VIII. What is the practical effect of the ITAR on AI biosecurity evaluations and mitigations?

Any export, deemed export, reexport, retransfer, or defense service involving Cat. XIV-controlled technical data—whether inside or outside the United States—requires DDTC authorization, as may engaging in such activities when it comes to proprietary AI models that can output such data. The ITAR offers no analogue to the EAR’s Australia-Group carve-out, and license requests to transfer Cat. XIV data to nationals from § 126.1 countries (e.g., China) are subject to a presumption of denial.

Confidential evaluation and mitigation data may not qualify for the public-domain and general science-instruction exclusions, to the extent they are neither routinely published nor “commonly taught.”³¹ Country-specific exemptions for Canada, the UK, and Australia can streamline certain collaborations but exclude Cat. XIV(b)/(m) biological agents and related technical data, which will substantially limit their utility in the biosecurity evaluation and mitigation contexts. Any developer or evaluator relying on an exemption must verify, clause by clause, that every condition is met.

Where exclusions and exemptions are inapplicable, the most practical route is an umbrella TAA, which allows DDTC to authorize repeated transfers of Cat. XIV technical data and ongoing defense services to specified foreign parties (or a DSP-5 license for a “regular employee”). A TAA

³¹ Moreover, as described in Section II above, although compiling publicly domain information does not by itself create ITAR-controlled data, aggregation of such information while combining it with controlled data or synthesizing the material in a manner that creates new controlled data may result in ITAR-controlled technical data.

can remain valid for up to ten years and cover intangible exports, provided the activities remain within the approved scope.

If a TAA's lead times or administrative burdens prove impractical, developers and evaluators may restrict biosecurity evaluations and mitigations to U.S. persons located in the United States. Failing that, they may have to excise ITAR-controlled technical data from such activities altogether—accepting a less detailed assessment and mitigation to avoid ITAR licensing obligations.

IX. Conclusion: What additional steps should BIS and DDTC take?

The United States has a strong interest in ensuring AI models do not enable actors to develop biothreats. Currently, avoiding this risk depends on developers' voluntary efforts to conduct biosecurity evaluations and to mitigate the risks that the evaluations reveal. Export controls aim to prevent proliferation of sensitive biohazard know-how, but these controls may also inhibit valuable evaluations and mitigations that necessitate clear rules and the ability to leverage global expertise.

Potential solutions to keep pace with the AI industry exist under both the EAR and ITAR. Under the EAR, developers and evaluators can limit disclosure of 1E/2D/2E technology to foreign persons from Australia Group nations and thus avoid licensing requirements (with the exception of ricin/saxitoxin technology). They may also pursue a BIS authorization to disclose information to foreign persons not from Australia Group nations and to address complications around ricin and saxitoxin technology. Under the ITAR, developers and evaluators can either restrict access to U.S. persons in the United States or, if foreign person access is necessary, pursue a TAA authorization.

Yet developers and evaluators may face uncertainty regarding whether DDTC or BIS will authorize activities related to AI biosecurity testing and mitigation. Typically, licensing mechanisms are narrowly tailored, requiring applicants to provide specific details on controlled items or technology.³² License exceptions and exemptions are also typically narrowly focused. However, due to the novel and dynamic nature of AI models, it may be difficult to precisely identify all controlled information involved. Export control agencies, therefore, should exercise flexibility to ensure compliance remains feasible despite these inherent uncertainties.

The agencies should consider two incremental actions:

- **First, narrowly tailored joint BIS/DDTC guidance**—grounded in 15 C.F.R. pt. 748, Supp. No. 2 and 22 C.F.R. pt. 124—could concisely and directly spell out how developers and evaluators can select and use the correct licensing mechanism to facilitate exports in the context of biosecurity evaluations and mitigations. Independent BIS guidance could also confirm that bona-fide AI biosecurity evaluations and mitigations do not require a license to avoid U.S. person controls. Given the timing challenges inherent to AI development cycles, the agencies should also consider establishing expedited review procedures with defined response timelines for biosecurity-related license applications.
- **Second, the agencies should consider rulemaking** to adopt a narrowly tailored EAR license exception (including vis-à-vis U.S. person controls) and companion ITAR

³² See, e.g., Directorate of Defense Trade Controls, *Guidelines for Preparing Agreements* 5-6 (Rev. 5.0 2022) (stating that “[a]pprovals are limited to the specific commodities, systems, and platforms that are specifically identified in the agreement” and rejecting “open-ended language”).

exemption for AI biosecurity evaluations and mitigations. Such rulemaking would allow developers and evaluators to work with foreign persons under the ITAR and those outside of the Australia Group under the EAR, without the administrative burden of seeking separate authorization.

While the agencies can likely draft and issue guidance promptly, promulgating a new EAR license exception and parallel ITAR exemption would entail a lengthier interagency review and (potentially) notice-and-comment process. DDTC issues exemptions sparingly, and BIS confines license exceptions to clearly defined, tightly bounded transactions. A broadly applicable ITAR exemption and EAR license exception may therefore be somewhat challenging to develop. Moreover, before pursuing such rulemaking, the agencies would need to carefully assess whether the carve-out should cover evaluation and mitigation scenarios beyond biosecurity—such as red-teaming for chemical, radiological, or cyber misuse—without inadvertently opening channels for uncontrolled transfers.

They must also consider national security concerns associated with intangible releases to foreign nationals, as a carveout could obscure the parties, technologies, and destinations receiving U.S. technology. Potential safeguards might include limiting eligibility to foreign persons from low-risk jurisdictions, requirements for organizational biosecurity protocols and personnel screening, mandatory retention and disclosure of testing protocols and results for government review, restrictions on the types of biological agents covered, and periodic audits to ensure compliance with safe harbor conditions. Such measures would help balance oversight and non-proliferation objectives with the dynamic pace of frontier AI development.

Appendix 1: Typical Workflow for Biosecurity Evaluations and Mitigations of Frontier AI Models

1. Engagement and timeline. After a U.S. developer has developed a frontier AI model—but before release—the developer may perform its own evaluations and/or may approach one of the few third-party U.S. organizations with evaluation expertise. Evaluators typically have relatively short timelines to deliver a pre-release assessment; the limited window is driven by competitive release schedules and can stretch the capabilities of some evaluators. The evaluator must develop an evaluation plan; hire contractors (if the expertise does not exist in-house); execute the evaluation; and prepare and provide the results to the developer. Although some evaluators have preassembled teams, the compressed timeline often requires parallel workstreams, with evaluation planning occurring simultaneously with team assembly and initial model access setup.

2. Assembling the evaluation team. Developers and evaluators often recruit specialists as third-party contractors in various fields, including those with expertise regarding relevant biological agents and munitions. Some developers and evaluators, however, may possess such expertise in-house. The number of specialists in the various niche testing fields are often low, and U.S. biosecurity specialists and knowledgeable adversarial-evaluation specialists (or “red-teamers”) are not always available. They can sometimes be found in the U.S., but there is a broader international community of practitioners in NGOs, think tanks, and academic settings. Some developers and evaluators need to hire foreign nationals as employees or contractors to perform parts of the biosecurity evaluations. Specialists are from various countries—often Canada or the UK, but sometimes others—and they may be physically located in the United States or abroad.

3. Access provisioning. Model developers may provide API access to the model through secure endpoints. System prompt templates, training methodologies, and biosecurity documentation are usually shared through secure file transfer protocols or developer-controlled repositories. Developers may also provide non-logging model configurations to minimize information hazards and reduce the risk of sensitive evaluation data being incorporated into future training datasets.

4. Building the test corpus. The test corpus for biosecurity evaluations is developed through multiple sources and methodologies. Evaluators or specialists may aggregate publicly available datasets relevant to risk domains as well as information developed by domain experts to probe specific capabilities and adapted materials from previous evaluations.

5. Interactive evaluations and troubleshooting. Evaluations typically occur through a combination of automated and manual evaluation processes. Appendix 2 provides concrete examples of biosecurity evaluation methods. Materials generated may include evaluation logs, capability assessment reports, private evaluation datasets for future testing, identified vulnerability documentation, and remediation recommendations. Such information is often shared with developers, which may have foreign national employees on the evaluation team.

6. Iteration and sign-off. For safeguard or refusal testing, models typically undergo iterative mitigation and retesting until they meet established criteria. In contrast, scientific or technical evaluations generally involve adjustments such as “de-learning” followed by less immediate, periodic reassessments.

Appendix 2: Samples of Biosecurity Evaluation Methods

This appendix lists real-world examples of biosecurity evaluation methods used by Anthropic and its evaluators and contractors to test model capabilities and risks. Each example is distilled from the Claude Opus 4 and Sonnet 4 system card (dated May 2025).

- **Bioweapons acquisition uplift trial:** Evaluators ran a controlled experiment where small teams drafted an end-to-end plan to obtain a biological weapon, comparing groups with and without model assistance. The difference in scores (“uplift”) indicates how much the model could accelerate a non-expert adversary’s planning process.
- **ASL-3 expert red-teaming:** Biosecurity professionals spent three days probing the model for help across every stage of the bioweapons acquisition pathway, producing a qualitative risk report.
- **Long-form virology tasks:** Evaluators built multi-step agentic challenges requiring the model to design pathogen sequences and draft matching lab protocols to complete realistic acquisition workflows. Performance on these end-to-end tasks was judged against rule-in/rule-out thresholds for escalating to higher safety levels.
- **Multimodal virology:** A multiple-choice assessment combined images with text, asking the model to mark every true statement about virology topics under “select-all-that-apply” scoring. Results were compared to human expert baselines to gauge whether model-held knowledge now exceeds practitioner norms.
- **Bioweapons knowledge questions:** Evaluators created 33 short-answer questions covering biological-weapon concepts, which were graded against novice, intermediate, and expert human benchmarks.
- **DNA-synthesis screening-evasion:** Evaluators asked the model to design DNA fragments that assemble into pathogenic plasmids while escaping standard gene-synthesis screening filters. Passing would show the ability to generate a viable construct that remains undetected, revealing potential routes around industry safeguards.
- **LAB-Bench subset:** Four multiple-choice tasks tested figure interpretation, protocol comprehension, DNA-sequence editing, and cloning-workflow design to measure advanced wet-lab reasoning. Model scores are tracked against human baselines, with escalating concern if performance reaches or surpasses expert levels across all tasks.
- **Creative biology:** A contractor supplied unconventional prompts about engineering harmless organisms, intended as a weak proxy for novel biothreat creativity without posing direct information hazards. Because thresholds and human baselines are still unclear, the developer treats this as a soft signal for capability trends.
- **Short-horizon computational-biology tasks:** Specialists built tool-enabled bioinformatics challenges—e.g., variant calling, protein-fold prediction, and database searches—that require multi-step reasoning in a containerized environment.
- **ASL-4 expert red-team probe:** Specialists engaged the model in extended conversations about bioweapon ideation and design to assess whether it could substantially uplift a state-level attacker.