
Responsible Advanced Artificial Intelligence Act

A BILL

To establish an administration that will oversee and regulate advanced general-purpose artificial intelligence systems.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SEC. 1. SHORT TITLE; TABLE OF CONTENTS.

(a) SHORT TITLE.—This Act may be cited as the “Responsible Advanced Artificial Intelligence Act of 2024”.

(b) TABLE OF CONTENTS.—The table of contents for this Act is as follows:

SEC. 1. SHORT TITLE; TABLE OF CONTENTS.
SEC. 2. SENSE OF CONGRESS.
SEC. 3. DEFINITIONS.
SEC. 4. FRONTIER ARTIFICIAL INTELLIGENCE SYSTEMS ADMINISTRATION.
SEC. 5. DEPUTY ADMINISTRATORS.
SEC. 6. RULEMAKING AUTHORITY.
SEC. 7. PRE-REGISTRATION FOR MEDIUM-CONCERN AI.
SEC. 8. PERMIT APPLICATIONS FOR FRONTIER AI.
SEC. 9. ADDITIONAL REQUIREMENTS FOR EXTREMELY HIGH-CONCERN AI.
SEC. 10. ADJUDICATION OF PERMIT APPLICATIONS.
SEC. 11. APPEALS OF PERMIT APPLICATIONS.
SEC. 12. HARDWARE MONITORING.
SEC. 13. ANALYSIS AND REPORTING.
SEC. 14. CIVIL LIABILITY.
SEC. 15. CRIMINAL LIABILITY.
SEC. 16. EMERGENCY POWERS.
SEC. 17. WHISTLEBLOWER PROTECTION.
SEC. 18. INTER-AGENCY COOPERATION.
SEC. 19. PREEMPTION.
SEC. 20. AUTHORIZATION OF FUNDING.
SEC. 21. SEVERABILITY.

SEC. 2. SENSE OF CONGRESS.

It is the sense of Congress that in recent years, artificial intelligence (AI) has rapidly grown more powerful. Computer scientists do not fully understand how AI systems work, nor do they know how to reliably control advanced AI systems. Without additional precautions, we cannot be confident that advanced AI systems will not develop bioweapons, launch

automated cyberattacks, manufacture armed drones, or otherwise cause a violent and dangerous catastrophe.

As leading AI developers themselves have acknowledged, private AI companies lack the right incentives to fully address this risk. Therefore, this Act will promote the safety of advanced AI development by creating a new administration that will monitor and regulate the use of AI on the frontier of modern capabilities. The administration will ensure that private developers of ‘frontier’ AI follow a set of security guidelines, will track concentrations of the semiconductors used for frontier AI, and will be prepared to rapidly intervene in case of an AI-related emergency.

SEC. 3. DEFINITIONS.

In this Act:

(a) **ADMINISTRATION.**—The term “Administration” means the Frontier Artificial Intelligence Systems Administration established under section 4 of this Act.

(b) **ADMINISTRATOR.**—The term “Administrator” means the Administrator of the Frontier Artificial Intelligence Systems Administration established under section 4 of this Act.

(c) **ARTIFICIAL INTELLIGENCE (AI).**—The terms “Artificial Intelligence” and “AI” each include the meanings assigned by Section 238(g) of P.L. 115-232 (the John S. McCain National Defense Authorization Act for FY2019) and by 8 P.L. 116-283; H.R. 6395, Division E, Section 5002(3) (the National Artificial Intelligence Initiative Act of 2020).

(d) **AI SYSTEM.**—The term “AI system” means a particular model, program, or tool within the field of AI.

(1) **COLLECTIONS OF SOFTWARE.**—For purposes of evaluating the tier into which an AI system should be classified, a collection of AI software may be treated as a single AI system if elements in the collection share a common purpose or design or are otherwise intended to function or do function as a coherent unit.

(2) **COLLECTIONS OF HARDWARE.**— Equipment may be considered a “collection of hardware” even if it has multiple owners or multiple locations as long as it is being used or made available for a common or coordinated purpose.

(e) **BOARD.**—The term “Board” means the Artificial Intelligence Appeals Board established under section 11 of this Act.

(f) **BOARD MEMBER.**—The term “Board Member” means a Board Member of the Artificial Intelligence Appeals Board established under section 11 of this Act.

(g) **COMPUTING POWER (COMPUTE).**—The term “computing power” and the term “compute” each refer to the processing power and other electronic resources used to train, validate, deploy, and run AI algorithms and models.

(h) **DEPLOYMENT.**—The term “deployment” means sharing or providing access to an AI system—

- (1) to the general public,
- (2) to an external audience,

(3) for the purpose of improving AI capabilities, except capabilities directly related to safety, alignment, robustness, or interpretability, or

(4) for any commercial purpose.

(i) DEVELOPMENT.—The term “development” includes research, engineering, testing, and the improvement of AI capabilities.

(j) FINAL TRAINING RUN.—The term “final training run” means all training that directly contributed to the determination of the model weights of an AI system, including pre-training, fine-tuning, reinforcement learning, and all other adjustments that influenced the same set of model weights.

(k) FLOP.—The term “FLOP” means a half-precision (16-bit) floating point operation, which is a measure of compute. This is the number of operations multiplied by the bitlength and divided by 16. The term “FLOPs” is the plural of “FLOP.” The term “FLOP/s” means floating point operations per second.

(l) FRONTIER AI SYSTEM.—The term “frontier AI system” means any AI system classified under this Act as “high-concern” or “extremely high-concern.”

(m) FRONTIER AI LAB.—The term “frontier AI lab” means any entity that meets at least one of the following criteria:

(1) It owns or controls a frontier hardware cluster.

(2) It has trained or is training a frontier AI System.

(3) It possesses (i.e., has control sufficient to operate) model weights that resulted from the training of a frontier AI System.

(n) FRONTIER HARDWARE CLUSTER.—The term “frontier hardware cluster” means a collection of hardware with a combined throughput of at least 10^{18} FLOP/s. Equipment may be considered a “collection of hardware” even if it has multiple owners or multiple locations as long as it is being used or made available for a common or coordinated purpose.

(o) HIGH-PERFORMANCE AI HARDWARE — The term “high-performance AI hardware” means any integrated circuit covered by ECCN 3A090.b in the October 25, 2023 Advanced Computing Chips Rule (AC/S IFR), 88 FR 73458, to wit, any single integrated circuit with (1) a total processing performance of 4800 or more, or (2) a total processing performance of 1600 or more and a performance density of 5.92 or more.

(p) IMMEDIATE FAMILY MEMBER.—The term “Immediate Family Member” means—

(1) a spouse or domestic partner, parent, grandparent, sibling, or child of the individual, including step, in-law, and adoptive relationships;

(2) any person to whom the individual stands in loco parentis; and

(3) any other person living in the household of the individual.

(q) MAJOR SECURITY RISK.—The term “major security risk” includes—

(1) risks that credibly threaten to substantially damage America’s public safety, essential infrastructure, or national security;

(2) global catastrophic and existential threats, as defined by the Global Catastrophic Risk Management Act of 2022; and

(3) risks that AI systems will establish self-replicating autonomous agents or otherwise permanently escape from human control.

(r) MODEL WEIGHTS.—The term “model weights” means the collection of parameters that transform input data into output data within some machine learning models, including neural networks.

(s) PRE-TRAINING.—The terms “pre-training” and “initial training” mean the process of training an AI system for the first time, giving the AI system a general suite of capabilities that may later be expanded through fine-tuning.

(t) SENIOR MANAGER.—The term “senior manager” includes:

(1) the Administrator,

(2) the Deputy Administrators,

(3) the Board Members,

(4) any individual whose duties would ordinarily be commensurate with a position in the Senior Executive Service, and

(5) any individual who is classified at the GS-13 level or above and who has any significant responsibility for determining policy or managing other professionals.

(u) TIERS OF CONCERN.—The terms “Tier 1,” “Tier 2,” “Tier 3,” and “Tier 4” have the meanings set forth below. A “higher tier” refers to a tier with a larger number and represents a more severe concern about major security risks; a “lower tier” refers to a tier with a smaller number and represents a less severe concern about major security risks. The Administrator may modify any or all of these thresholds as set forth in Sections 4 through 6.

(1) LOW-CONCERN AI SYSTEM (TIER 1).—The terms “low-concern AI system” and “Tier 1” mean AI systems that do not have any capabilities that are likely to pose major security risks. Initially, an AI system shall be deemed low-concern if it used less than 10^{24} FLOP during its final training run.

(2) MEDIUM-CONCERN AI SYSTEM (TIER 2).—The terms “medium-concern AI system” and “Tier 2” mean AI systems that have a small chance of acquiring at least one capability that could pose major security risks. For example, if they are somewhat more powerful or somewhat less well-controlled than expected, such systems might substantially accelerate the development of threats such as bioweapons, cyberattacks, and fully autonomous artificial agents. Initially, an AI system shall be deemed medium-concern if it used at least 10^{24} FLOP during its final training run and it does not meet the criteria for any higher tier.

(3) HIGH-CONCERN AI SYSTEM (TIER 3).—The terms “high-concern AI system” and “Tier 3” mean AI systems that have at least one capability that could pose major security risks, or that have capabilities that are at or very near the frontier of AI development, and as such pose important threats that are not yet fully understood.

(A) INITIAL CRITERIA.—Initially, an AI system shall be deemed high-concern if it used at least 10^{26} FLOP during its final training run or was ranked as one of the top five most capable AI systems in the world based on a set of benchmarks selected by the Deputy Administrator for Standards, as described in Section 5.

(B) EXCEPTIONS.

(i) An AI system that has a narrow purpose or is otherwise very unlikely to cause major security risks shall instead be classified as medium-concern, pursuant to the fast-track exemption procedures described in Section 8.

(ii) An AI system shall instead be classified as extremely-high-concern if it qualifies as Tier 4 based on the criteria in paragraph (4).

(4) EXTREMELY HIGH-CONCERN AI SYSTEM (Tier 4).—The terms “extremely high-concern AI system” and “Tier 4” mean AI systems whose training, development, or deployment are highly likely to pose major security risks. Initially, an AI system shall be deemed extremely high-concern based on any of the following threats—

(A) the AI system has or could easily develop the ability to significantly assist with the development, manufacture, or deployment of biological, chemical, radiological, or nuclear weapons;

(B) the AI system has or could easily develop the ability to autonomously spread, replicate in an uncontrolled manner, or take over external computers;

(C) the AI system has or could easily develop the ability to accelerate scientific research to such a degree as to undermine national security or destabilize the global balance of power;

(D) some or all of the AI system’s capabilities significantly exceed normal human levels of performance on one or more tasks relevant to major security risks; or

(E) the AI system otherwise poses significant existential or global catastrophic risks.

(v) TRAINING.—The term “Training” means the process of fitting model weights and biases to a machine learning algorithm so it can build a representation of the relationship between data features and a target label or among the features themselves. This process teaches an AI system to perceive, interpret, and learn from data so it can be capable of reaching conclusions that are based on that data. Training includes both pre-training and fine-tuning.

SEC. 4. FRONTIER ARTIFICIAL INTELLIGENCE SYSTEMS ADMINISTRATION.

(a) ESTABLISHMENT.—There is established a Federal administration, to be known as the “Frontier Artificial Intelligence Systems Administration”, which shall—

(1) be constituted as provided in this Act; and

(2) execute and enforce the provisions of this Act.

(b) ADMINISTRATOR.—

(1) HOW APPOINTED.—The Administrator shall be appointed by the President with the advice and consent of the Senate, on the basis of the Administrator’s demonstrated leadership or management experience at the intersection of security and advanced technology, such as a background in cybersecurity, biosecurity, or existential risks from other advanced technologies.

(2) HOW REMOVED.—The Administrator shall serve at the pleasure of the President, unless removed by impeachment.

(3) RESPONSIBILITIES.—The Administrator has ultimate responsibility for exercising all powers and responsibilities pursuant to this Act.

(4) COMPENSATION.—The Administrator is entitled to be compensated as a Level III Executive under 5 USC § 5314.

(5) DELEGATION OF AUTHORITY.—The Administrator shall delegate responsibilities to the Deputy Administrators as set forth in Section 5. Each Deputy Administrator shall perform their delegated responsibilities subject to the lawful instructions of the Administrator. The Administrator may remove a Deputy Administrator for cause; if the Administrator does so, then the Administrator shall appoint a replacement within 30 days. The Administrator may delegate (and subsequently resume) other responsibilities to any person employed by the Administration.

(b) CONFLICTS OF INTEREST.—No person may be appointed or serve as a senior manager under this Act who has any conflict of interest.

(1) HOW DEFINED.—A conflict of interest includes—

(A) owning any stocks, bonds, options, or other interest in any company that develops, sells, or promotes artificial intelligence, except in so far as such interest is wholly contained in a blind trust or public mutual fund.

(B) serving on the board of directors, board of trustees, or similar advisory board for any frontier AI lab;

(C) lobbying on behalf of any frontier AI lab at any time within three years of the first day of the senior manager’s service under this Act;

(D) working as an employee or contractor (other than as a lobbyist) of any frontier AI lab at any time within three years of the first day of the senior manager’s service under this Act; or

(E) having an immediate family member who meets any of the criteria in subparagraphs (A) through (D) above.

(2) WAIVERS.—The Administrator may waive a conflict of interest under subparagraph (1)(D) or (1)(E) based on a written and dated memorandum that is personally signed by the Administrator, finding that both (1) the senior manager’s skills cannot be adequately replaced despite the conduct of a diligent search, and (2) the conflict does not pose a significant threat to the integrity of the Administration. The Administrator may not waive a conflict of interest as to the Administrator’s own service.

(3) HOW RESOLVED.—Upon the discovery that a senior manager has served under this Act despite the existence of a significant un-waived conflict of interest, that senior

manager shall immediately resign or be discharged from service, and then the Administrator shall promptly review all significant actions taken by that senior manager and may endorse, amend, or reverse each such action based on the minimum requirements of constitutional due process, notwithstanding any other procedural requirement. However, the fact that the senior manager was not eligible to serve shall not otherwise impair the validity of the acts taken by the senior manager during their service.

(4) LIMITATIONS ON SUBSEQUENT EMPLOYMENT.

(A) PROHIBITED EMPLOYMENT.—Any person who has served as a senior manager under this Act shall not allow themselves to acquire any conflict of interest under subparagraphs (f)(1)(A), (f)(1)(C), or (f)(1)(D) for a period of three years following the last day of their service. Similarly, no such person may lobby the Administration for a period of three years following the last day of their service.

(B) UNREASONABLE COMPENSATION.—Any person who has served as a senior manager under this Act and who acquires a conflict of interest under subparagraph (f)(1)(B) within three years following the last day of their service shall refuse any payment of excessive or unreasonably high compensation. Likewise, any frontier AI lab hiring such a person or their immediate family member shall take care to pay only reasonable compensation that represents the fair value of the person’s skill and effort, and that does not suggest a *quid pro quo* for political favors.

(c) EMPLOYEES.—

(1) IN GENERAL.—The Administrator may, subject to the civil service laws and the Classification Act of 1949, as amended, hire such employees as are useful in the exercise of the Administration’s functions.

(2) ADMINISTRATIVE ASSISTANTS.—Without regard to the civil service laws, the Administrator and each Deputy Administrator may each appoint one administrative assistant, who shall perform such duties as that senior manager shall direct. Each such administrative assistant shall be entitled to compensation based on the GS-9 pay scale.

(3) PRIORITY POSITIONS.—The Administrator may designate priority positions.

(A) LIMIT ON QUANTITY.—The Administrator may not employ more than 400 full-time equivalent personnel under priority positions at any one time.

(B) BASIS FOR DESIGNATION.—To be designated as a priority position, a job role must require skills that are rare, advanced, highly in demand by the private sector, or otherwise difficult for the Administration to acquire.

(C) DIRECT HIRING AUTHORITY.—After taking into consideration the availability of preference eligibles for the position (as defined by 5 USC 2108), the Administrator may directly hire individuals for priority positions, without regard to the provisions of any other law relating to the appointment, number, classification, or compensation of employees.

(D) TECHNICAL RECRUITING.—When making appointments under this paragraph, the Administrator shall take care to accurately describe (i) the technical skills needed for each position, (ii) the software and other tools that will be used in each position,

and (iii) the job duties of each position. The Administrator shall consult technical experts as necessary in order to make these descriptions accurate.

(E) BONUS PAY.—The Administrator shall designate a rate of basic pay for each priority position that is 150 percent of the rate that would ordinarily be applied to that position's classification. The Administrator may refer classification decisions to the Office of Personnel Management. Locality pay adjustments and similar benefits for priority positions shall be calculated based on this increased rate of basic pay.

(4) HONORS PROGRAM.—The Administrator is encouraged and empowered to design and implement a recruiting program for talented entry-level professionals comparable to the Honors Program of the Department of Justice.

(d) VOLUNTEERS.—The Administration may accept volunteers.

(1) Such volunteers shall not be barred from service by part III of title 5, United States Code or section 1342 of title 31, United States Code.

(2) Any individual who provides voluntary services under this subsection or who provides goods in connection with such voluntary services shall not by reason of such voluntary service be considered a Federal or special government employee.

(3) Any individual who provides voluntary service under this subsection shall first sign a waiver indicating that their voluntary service is provided without any hope of reimbursement, and expressly waiving any claim for payment for said service.

(4) No person may volunteer for the Administration while that person is employed or receiving any compensation (other than minor gifts that would be permitted under the Congressional Gift Rule at House Rule 25, clause 5) from any company that develops, sells, or promotes artificial intelligence.

(e) DONATIONS.—The Administration may accept donations (including gifts and bequests) to be used in the furtherance of its functions.

(1) LIMITATION ON CONDITIONS.—Normally, the Administration may not accept a donation that has any condition attached to it. However, the Administration may accept a donation that has one or more of the following conditions attached to it:

(A) A condition that directs the donation toward a division, bureau, or program within the Administration, e.g., toward monitoring, or toward standards, or toward enforcement.

(B) A condition that directs that such portion of the donation that is not spent by a particular date be returned to the donor.

(2) NO REAL OR APPARENT CONFLICTS OF INTEREST.—The Administration may not accept a donation if that donation would create a real or apparent conflict of interest.

(3) TAXES—For the purpose of Federal law on income taxes, estate taxes, and gift taxes, property or services accepted under this subsection shall be deemed to be a gift, bequest, or devise to the United States.

(4) MANAGEMENT OF GIFTS.—Insurance, interest, accounting, and similar management of any gifts received under this subsection shall be handled using the same

procedures specified under 42 U.S. Code § 238, except that the Administrator shall perform any duties assigned by that section to the Secretary of Health and Human Services or to the Surgeon General, and any references to the Public Health Service shall instead refer to the Administration.

(f) PAPERWORK REDUCTION ACT.—The Administration is not required to wait for prior approval under the Paperwork Reduction Act before collecting information from persons who are or may be frontier AI labs. However, the Administration must follow any lawful instructions issued by OMB directly to the Administration with respect to the forms and frequency of paperwork.

SEC. 5. DEPUTY ADMINISTRATORS.

(a) HOW APPOINTED.—The Administrator shall appoint each of the Deputy Administrators named in this Section within 60 days after the Administrator is appointed. If a Deputy Administrator position under this Section becomes vacant for any reason, the Administrator shall appoint a replacement within 45 days. Each Deputy Administrator shall be appointed based on that Deputy Administrator’s demonstrated skill and experience in the field of computer science or in the areas for which that Deputy Administrator is responsible.

(1) The Deputy Administrator for Public Interest shall also be appointed based on their integrity, high moral character, and independence from relevant commercial interests.

(2) The Administrator may leave the position of Deputy Administrator for Grants Management vacant during any year in which Congress has appropriated less than \$1 million in grants to be disbursed under this Act.

(b) COMPENSATION.—Each Deputy Administrator is entitled to be compensated as a Level IV Executive under 5 USC § 5315.

(c) DEPUTY ADMINISTRATOR FOR MONITORING.—The Deputy Administrator for Monitoring shall be responsible for supervising and directing the progress of hardware monitoring and reporting under Section 8 of this Act. In particular, the Deputy Administrator for Monitoring shall develop and maintain an awareness of the physical locations and ownership of high-performance AI hardware, and shall organize efforts to detect and identify any high-performance AI hardware that has not been properly accounted for. In addition, the Deputy Administrator for Monitoring is the first assistant to the Administrator pursuant to the Federal Vacancies Reform Act of 1998.

(d) DEPUTY ADMINISTRATOR FOR STANDARDS.—The Deputy Administrator for Standards shall be responsible for supervising and directing the progress of rulemaking under Sections 6 and 7 of this Act. In particular, the Deputy Administrator for Standards shall—

(1) issue, evaluate, and regularly update the rules governing applications for frontier AI permits;

(2) take care to ensure that the standards for issuing a frontier AI permit are strict enough to adequately protect against major security risks;

(3) regularly update the thresholds for the tiers of concern;

(4) within 60 days after being appointed, select and publish in the Federal Register a set of one or more benchmarks to be used in the definition of “high-concern AI systems”

from among those benchmarks that are commonly used to quantify the performance of state-of-the-art foundation models, that are established by industry best practices, or that are endorsed by relevant standard-setting organizations; and

(5) within 60 days after being appointed, establish a weighting system for the benchmarks in paragraph (4) above, such as averaging the scores on each benchmark.

(e) DEPUTY ADMINISTRATOR FOR ENFORCEMENT.—The Deputy Administrator for Enforcement shall be responsible for investigating, prioritizing, and prosecuting violations of this Act with the goal of deterring actual and potential violators from taking actions that pose major security risks. In so doing, the Deputy Administrator for Enforcement may directly pursue civil and criminal cases, and the Deputy Administrator for Enforcement may refer civil and criminal cases to the Department of Justice.

(f) DEPUTY ADMINISTRATOR FOR ALGORITHMS.—The Deputy Administrator for Algorithms shall be responsible for observing and evaluating the rate of progress in the algorithmic efficiency of artificial intelligence, for making recommendations to the Administrator on how and when to update the technical definitions in this Act (including the definition of frontier AI) so as to compensate for such progress and prevent an increase in major security risks, and for ensuring that such updates are promulgated and become final in a timely fashion.

(g) DEPUTY ADMINISTRATOR FOR GRANTS MANAGEMENT.—The Deputy Administrator for Grants Management shall be responsible for designing, awarding, obligating, managing, and evaluating the outcomes from grants issued for the following purposes:

(i) PUBLIC COMPUTE BANK.—Acquiring compute and suitable supporting facilities and then lending those resources out in the public interest for free or at a discounted cost so as to allow academics, researchers, advocates, and non-profit entities to test, evaluate, and explore the implications of AI systems.

(ii) HARDWARE SAFETY RESEARCH.—Funding research, development, and prototypes of hardware safety features for AI systems, especially on-chip features that enhance the transparency or verifiability of high-performance AI hardware or that render such hardware easier to remotely monitor or remotely disable.

(iii) IMPROVED EVALUATION TECHNIQUES.—Funding research, development, and prototypes of improved evaluations for AI systems, especially evaluations that capture information about the safety or alignment of an AI system and evaluations that can be applied automatically, objectively, quickly, or at scale.

(iv) VOLUNTARY AUDITORS FOR SMALL BUSINESS.—Funding auditing, red-teaming, or similar safety evaluations or protocols for small businesses or entrepreneurs, especially when such businesses would otherwise be at a competitive disadvantage against larger or more established providers of AI if they attempted to match the safety features of those larger providers.

(h) DEPUTY ADMINISTRATOR FOR PUBLIC INTEREST.—The Deputy Administrator for Public Interest shall be responsible for assessing the impact of frontier AI systems on the rights and interests of the American public, consulting with non-governmental organizations (especially organizations that evaluate AI models) as to the risks posed by frontier AI systems, warning appropriate officials and the public about dangers in the field of AI, nominating

members for appointment by the Administrator to the Appeals Board, investigating and reporting on frontier AI systems of special concern, and advocating for the denial of permits for AI projects that would increase major security risks.

(1) **ADDITIONAL POWERS.**

(A) The Deputy Administrator for Public Interest may initiate complaints or proceedings before the Administration or the Board on behalf of itself or the American public in relation to the permitting of frontier AI.

(2) The Deputy Administrator for Public Interest may appear or intervene, as a party or otherwise, as a matter of right before the Administration, the Board, or any Court reviewing any action done pursuant to this Act. In any such appearance, the standing of the Deputy Administrator shall include standing to advocate any position on any matter involving the permitting of frontier AI systems or involving rules or procedures of the Administration affecting the American public.

(3) The Deputy Administrator for Public Interest is entitled to access to records available in a proceeding before the Administration or the Board and to obtain discovery of any nonprivileged matter that is relevant to its functions, subject to confidentiality requirements.

(2) **ADDITIONAL PROTECTIONS.**—Neither the Administrator nor the Acting Administrator may delay, hinder, prevent, or prohibit the Deputy Administrator for Public Interest from—

(A) initiating, carrying out, or completing any audit or investigation;

(B) issuing any subpoena during the course of any audit or investigation;

(C) filing any lawsuit or maintaining any argument or position pursuant to such lawsuit;

(D) advocating for the rejection or modification of any frontier AI permit; or

(E) selecting the content, timing, and audience for any report to the President, the Congress, or the general public.

(i) **DEPUTY ADMINISTRATOR FOR EMERGENCY PLANNING.**—The Deputy Administrator for Emergency Planning shall be responsible for preparing and planning for AI-related emergencies, and, if necessary, implementing any policies ordered by the President or by the Administrator pursuant to the emergency powers authorized by section 11 of this Act.

SEC. 6. RULEMAKING AUTHORITY.

(a) **IN GENERAL.**—Except as otherwise modified by this section, the Administrator shall have full power to promulgate rules to carry out this Act in accordance with section 553 of title 5, United States Code. This includes the power to update or modify any of the technical definitions in this Act (including the definition of “frontier AI” and the definitions of “tiers of concern”) to ensure that these definitions will continue to adequately protect against major security risks despite changes in the technical landscape such as improvements in algorithmic efficiency. However, neither these definitions nor any rule promulgated under this Act may be altered by the Administrator so as to be more permissive of frontier AI development unless

the Administrator first makes findings supported by clear and convincing evidence that such alterations will not significantly increase major security risks.

(b) ABBREVIATED NOTICE FOR CHANGES TO PERMITTING STANDARDS.—The Administrator may publish an advance notice of proposed rulemaking in the Federal Register no less than 10 days before the proposed rule goes into effect.

(1) CONDITIONS.—The Administrator may only employ this abbreviated notice when all of the following conditions are met:

(A) The rule's only effects are to adjust one or more requirements for obtaining or maintaining one or more frontier AI system permits under Section 8 based on changes in available technology or software.

(B) The rule is no more than 500 words long.

(C) The rule is accompanied by and supported by the Administrator's written finding that the adjustments are needed to maintain an appropriate level of public safety in light of specific changes in the AI technology or AI software that is becoming available.

(2) UPDATES.—If, after the 10 day period of abbreviated notice, the Administrator revises the proposed rule based on comments received during the notice-and-comment process, and the revised rule still addresses substantially the same topic, and the revised rule still meets the conditions listed in paragraph (1), then the Administrator may re-publish the revised rule in the Federal Register no less than 5 days before the revised proposed rule goes into effect. After re-publication, the Administrator is not required to respond to comments that do not relate to the new revisions.

(c) EXTENDED NOTICE FOR GENERAL BEST PRACTICES.—The Administrator may issue one or more rules designed to reduce major security risks from AI. Each such rule shall concern best practices for the research, development, training, or deployment of AI. If the Administrator issues such a rule, and the rule is not directly related to the frontier AI permitting regime under Section 8, then the Administrator shall do all of the following:

(1) ANPRM.—The Administrator shall issue an Advance Notice of Proposed Rulemaking no less than 90 days before the proposed rule would go into effect.

(2) LISTENING SESSIONS.—The Administrator shall hold at least 2 listening sessions at which a broad cross-section of stakeholders are invited to discuss the proposed rule with the Administration, each of which shall be completed prior to the issuance of the Notice of Proposed Rulemaking.

(d) PETITION FOR RULEMAKING.—Pursuant to 5 U.S.C. §553(e), any person may petition the Administration to issue, amend, or repeal a rule on any topic within the Administration's authority. Within 60 days after receiving each such petition, the Administrator shall reply with a definite written statement as to the Administrator's intentions with respect to that petition. If the Administrator's intentions include future action, then the Administrator shall specify the date by which that action will be taken. A person who submits such a petition is entitled to a reply within 60 days. The Administrator's statement (or lack thereof) is subject to judicial review by the Federal District Court for the District of Columbia.

(e) PETITION FOR ACTION.—If this Act requires the Administrator to perform any action by a certain date, and 30 days have passed since the expiration of that date, then any person may petition the Administrator to perform the action. Within 15 days after receiving each such petition, the Administrator shall reply with a definite written statement as to the Administrator’s intentions with respect to that petition. If the Administrator’s intentions include future action, then the Administrator shall specify the date by which that action will be taken. A person who submits such a petition is entitled to a reply within 15 days. The Administrator’s statement (or lack thereof) is subject to judicial review by the Federal District Court for the District of Columbia.

(f) CONGRESSIONAL REVIEW ACT.

(1) The Administrator may make a determination pursuant to 5 U.S.C. §801(c) that a rule issued by the Administrator should take effect without further delay because avoidance of such delay is necessary to reduce or contain a major security risk. If the Administrator makes such a determination and submits written notice of such determination to the Congress, then a rule that would not take effect by reason of 5 U.S.C. §801(a)(3) shall nevertheless take effect. The exercise of this authority shall have no effect on the procedures of 5 U.S.C. § 802 or on the effect of a joint Congressional resolution of disapproval.

(2) Because of the rapidly changing and highly sensitive technical landscape, a rule that appears superficially similar to a rule that has been disapproved by Congress may nevertheless be a substantially different rule. Therefore, a rule issued under this section that varies at least one material threshold or material consequence by at least 20% from a previously disapproved rule is not “substantially the same” under 5 U.S.C. § 802(b)(2).

(g) MAJOR QUESTIONS DOCTRINE.—It is the intent of Congress to delegate to the Administration the authority to mitigate the major security risks of advanced, general-purpose artificial intelligence using any and all of the methods described in this Act. The Administration is expected and encouraged to rapidly develop comparative expertise in the evaluation of such risks and in the evaluation of the adequacy of measures intended to mitigate these risks. The Administration is expressly authorized to make policy judgments regarding which safety measures are necessary in this regard. This Act shall be interpreted broadly, with the goal of ensuring that the Administration has the flexibility to adequately discharge its important responsibilities.

(h) NO EFFECT ON EMERGENCY POWERS.—Nothing in this section shall be construed to limit the emergency powers granted by Section 11.

(i) STANDARD FOR REVIEW.—In reviewing a rule promulgated under this Act that increases the strictness of any definition or scoring criterion related to frontier AI, a court may not weaken or set aside that rule unless there is clear and convincing evidence of at least one of the following—

- (1) doing so will not pose major security risks, or
- (2) the rule exceeded the Administrator’s authority.

SEC. 7. PRE-REGISTRATION FOR MEDIUM-CONCERN AI.

(a) **SCOPE OF REQUIREMENT.**—Each person who trains any medium-concern AI system shall, before conducting any substantial part of that training, pre-register their training plan with the Administration. The pre-registration notice shall include all of the following information:

- (1) The name and contact information of the person responsible for the training.
- (2) The maximum amount of compute to be used during the largest training run, including all fine-tuning, reinforcement learning, and all other modifications that directly influence the final model weights that result from the training.
- (3) The general purpose of the AI system being trained.
- (4) The final scores of the AI system on the benchmarks selected by the Deputy Administrator for Standards.
- (5) The principal address at which the AI system is being trained, or, if the AI system is trained using cloud computing resources, the name of the principal cloud computing provider used for the training.

(b) **CONTINUAL TESTING.**—At the conclusion of each training run, or at least once each month (whichever is more frequent), each person who trains any medium-concern AI system shall test the AI system against each of the benchmarks listed in Section 3(v)(3)(a)(ii).

(c) **DISCOVERY OF HIGH-CONCERN PERFORMANCE.**—If, at any time during the training of a medium-concern AI system, that system’s scores on the benchmarks listed in Section 3(v)(3)(a)(ii) exceed an average of 80%, then the trainer shall immediately stop work on that AI system. No person may resume such work unless and until they have received a valid training permit for a high-concern AI system. To the maximum extent possible, access to that AI system must be cancelled unless and until the system receives a valid deployment permit for a high-concern AI system.

SEC. 8. PERMIT APPLICATIONS FOR FRONTIER AI.

(a) **SCOPE OF REQUIREMENT.**—Each person who meets one or more of the criteria listed below must qualify for, apply for, receive, and maintain the appropriate type(s) of frontier AI permit(s) from the Administration:

(1) **HARDWARE PERMIT.**—The person owns, imports, leases, rents, knowingly possesses, or knowingly uses a frontier hardware cluster.

(2) **TRAINING PERMIT.**—The person has trained, is training, or intends to train a frontier AI system. For components of the definition of “frontier AI system” that refer to evaluations that can only be performed after training, the person shall make a prediction of how well the AI system will perform on each such evaluation and include those predictions in the person’s application for a training permit. The AI system’s actual performance on those evaluations shall be reported in each of the person’s applications for model weights permits and deployment permits for that AI system, and the Administrator shall evaluate those applications in part based on how well the AI system’s actual performance corresponds to its predicted performance.

(3) **MODEL WEIGHTS PERMIT.**—The person owns, knowingly possesses, knowingly accesses, or knowingly exploits the model weights for a frontier AI system, except that—

(A) a consumer or end user who interacts with the model weights for a frontier AI system only by making use of a product or service that relies on those model weights does not require a permit; and

(B) a person who trains an AI system without intending to develop the capabilities of a frontier AI system who nevertheless accidentally creates a frontier AI system may retain up to 2 copies of the resulting model weights on secure servers pending a final decision on an application for a model weights permit, provided that such a person shall (i) immediately notify the Administration that they have accidentally created a frontier AI system, (ii) immediately cease all training and development of that AI system unless and until their model weights permit is approved, and (iii) apply for a model weights permit within 30 days after creating the frontier AI system.

(4) DEPLOYMENT PERMIT.—The person shares, publishes, distributes, rents, sells, or otherwise provides access to a frontier AI system to anyone other than—

(A) a dedicated safety tester;

(B) a government official who has lawfully requested such access; or

(C) a person performing training on that frontier AI system pursuant to a valid training permit.

(b) BREADTH OF PERMITS.—Each permit under this section applies only to a specific item or activity, and not to all AI-related activities that a person might wish to conduct.

(1) BREADTH OF HARDWARE PERMIT.—A hardware permit authorizes possession only of one specific collection of hardware; if the person wishes to substantially change the collection or acquire another collection of hardware, then the person needs an additional permit.

(2) BREADTH OF TRAINING PERMIT.—A training permit authorizes training only of one specific AI system (including all of that system’s checkpoints); if the person wishes to add additional features to the AI system that were not included in the original training permit, or if the person wishes to train a new AI system, then the person needs an additional permit.

(3) BREADTH OF DEPLOYMENT PERMIT.—A deployment permit authorizes deployment only of one specific AI system (in its final version) to the set of users described in the permit; if the person wishes to substantially change the AI system being deployed, or to substantially change the user base, or to deploy a new AI system, then the person needs an additional permit.

(4) BREADTH OF MODEL WEIGHTS PERMIT.—A model weights permit authorizes possession only of one specific set of model weights; if the person wishes to acquire other model weights, then the person needs an additional permit.

(c) TIMING OF REQUIREMENT.

(1) GRANDFATHER CLAUSE.—A person who has already acquired hardware, trained an AI system, deployed an AI system, or acquired AI model weights as of the date this law is enacted may indefinitely continue owning or using that hardware, AI system, or set

of model weights unless such activities are otherwise prohibited by law; such a person is not required to apply for a permit.

(A) NARROW CONSTRUCTION.—This exemption is to be construed narrowly. For example, a person may not conduct a new training run on a frontier AI system without a permit simply because they conducted a similar training run before this law was passed.

(2) ONGOING ACTIVITIES.—A person who needs a permit under section 8(a) based on ongoing activities that began prior to the date when the Administrator first publishes application forms for the relevant permit type must apply for that permit within 60 days of when the form is published. If they do not do so, they must discontinue the activities within 60 days of when the form is published. If they do apply for a permit, they may continue the activities covered by the application while that application is being processed or appealed unless—

(A) the application is rejected by both the AILJs and the Appeals Board;

(B) the application is approved with conditions, and the person does not accept the conditions; or

(C) the application is returned for revisions, and the person does not submit the revisions within 30 days after the return date.

(3) NEW ACTIVITIES.—A person who needs a permit under section 8(a) based on activities they are conducting after the Administrator first publishes application forms for the relevant permit type must apply for that permit and receive the permit before beginning those activities.

(c) REVIEW OF THRESHOLDS.—No later than September 1st of each year, the Administrator shall review each of the thresholds in section 8(a), together with the relevant definitions in section 3, and determine whether each threshold and each definition remains adequate to defend against major security risks. If any threshold or definition has become inadequate, then the Administrator shall promptly promulgate rules to appropriately strengthen or tighten the threshold or definition. The Administrator is not required to review these thresholds or definitions during the same calendar year that the Act is enacted.

(d) APPLICATION PROCESS.—

(1) “FAST TRACK” EXEMPTION FORMS.—No later than four months following the enactment of this Act, the Administrator shall promulgate a “fast track” exemption form and a set of standards for evaluating that form.

(A) LENGTH OF FORM.—The fast-track exemption form shall be no more than two letter-sized pages in length and shall use standard font sizes.

(B) PURPOSE OF FORM.—The purpose of this form is to allow for the rapid review of AI systems that are extremely unlikely to pose major security risks. The developers of such systems should be promptly permitted to conduct their business.

(C) METHOD FOR EVALUATING FORM.—The use of rubrics and formal adjudication are not mandatory for fast track exemption forms; the Administrator may instead evaluate fast track exemption requests using any convenient method.

(D) WHO QUALIFIES FOR FAST TRACK.—An AI System that might meet the technical definition of a frontier AI system but that is narrow-purpose, well-understood, or otherwise extremely unlikely to pose major security risks should receive an exemption based on a “fast track” form, unless the AI system is integrated with a more dangerous AI system that does pose major security risks. Examples of AI systems that should usually be exempted through the use of the “fast track” form include—

- (i) self-driving cars;
- (ii) navigational systems;
- (iii) recommendation engines;
- (iv) fraud detection systems;
- (v) weather forecasting tools;
- (vi) tools for locating deposits of oil, gas, or minerals;
- (vii) AI systems designed to predict the demand, supply, price, cost, or transportation needs of products or services;
- (viii) search engines whose primary function is to suggest webpages; and
- (ix) AI systems whose function is substantially limited to generating still images, each of which typically contains no more than thirty words of text.

(2) INITIAL APPLICATION FORMS.—No later than six months following the enactment of this Act, the Administrator shall promulgate an initial application form for each of the four types of frontier AI permits (hardware, training, model weights, and deployment). Each type of permit shall have its own separate initial application form.

(3) RENEWAL APPLICATION FORMS.—No later than twelve months following the enactment of this Act, the Administrator shall promulgate a renewal application form for each of the four types of frontier AI permits (hardware, training, model weights, and deployment). Each type of permit shall have its own separate renewal application form. Each permit shall expire and require renewal 1 year after it is issued, unless the Administrator promulgates a rule varying this time period.

(4) UPDATES TO APPLICATION FORMS.—The Administrator may update application forms at any time. Any updated application forms shall be published on the Administration’s website. Any application submitted 60 or more days after the Administrator publishes an updated form shall use the updated form.

(5) CONTENT OF FORMS.—Each application form shall collect sufficient information for the Administrator to adequately evaluate whether a proposed activity related to frontier AI poses an unacceptable major security risk.

(6) APPLICATION FEES.—The Administrator may promulgate rules establishing an application fee to be paid by each applicant, which may vary based on the type of permit being applied for, the size and purpose of the entity requesting the permit, and whether the permit was granted.

(i) RESEARCH EXEMPTION.—An applicant creating an AI system for the primary purpose of conducting academic research on safety, fairness, transparency, equity, privacy, robustness, or reliability shall not be required to pay any application fee.

(ii) OPEN SOURCE EXEMPTION.—An applicant creating an AI system as a collaboration among volunteers who have committed to making any resulting products or services available to the public for free or at cost shall not be required to pay any application fee.

(iii) FAST TRACK EXEMPTION.—The Administrator may not charge a fee for the fast track exemption process.

(iv) SUPPORT FOR SMALL BUSINESS.—The Administrator shall take care that the amount and structure of any application fee does not disadvantage small businesses or entrepreneurs compared to large or established providers of AI. The Administrator shall set aside between 1% and 10% of any application fees collected in order to provide technical assistance and support to small businesses and entrepreneurs to help them complete applications.

(e) RUBRICS FOR APPLICATIONS.—No later than eight months following the enactment of this Act, the Administrator shall promulgate rubrics that explain how the Administrator will score each application, and what scores will be required in order for an application for a frontier AI permit to be approved.

(1) PRIORITIES.—These rubrics shall prioritize the need to mitigate major security risks. Subject to that restriction, the Administrator may assign any set of weights to the criteria in the rubrics and may choose any set of thresholds or scoring requirements for an application to be approved.

(2) MISSING CRITERIA.—In developing each rubric, the Administrator shall consider assigning scoring factors based on the criteria listed in the remainder of this section. For each such criterion that is not incorporated into a scoring factor for a relevant type of frontier AI permit, the Administrator shall issue a written statement in the instructions for the application for that type of frontier AI permit explaining why the criterion was not included.

(3) POWER TO REQUIRE PRECAUTIONS.—A rubric may encourage or require an applicant to adopt particular safety precautions as a condition of receiving or renewing a permit. These precautions may include any of the following:

(A) The applicant shall conduct a third-party evaluation or audit that covers specific topics, offers specific guarantees of the evaluator's independence, and makes specific assurances about the safety of the applicant's proposed activities.

(B) The applicant shall demonstrate, through penetration testing or otherwise, that the site(s) at which it will conduct activities under the permit are secure against specific types of hacking.

(C) The applicant shall not use more than a specific amount of compute for certain purposes, or at certain times, or at all.

(D) The applicant shall provide watermarks, labels, or other assurances that the products of its frontier AI systems will be traceable to the applicant.

(E) Other precautions that the Administrator may find useful or appropriate.

(4) MANDATORY INSURANCE REQUIREMENT.—Each rubric shall include a requirement that the applicant shall procure and place liability insurance, and shall specify the minimum coverage amount and minimum scope of coverage.

(5) UPDATES TO RUBRICS.—The Administrator may promulgate updates to these rubrics at any time, except that the Administrator may not promulgate an update that would weaken or loosen the rubrics without first making a written finding supported by clear and convincing evidence that the update will not significantly increase major security risks and publishing that finding and its support in the Federal Register at least 45 days before that update takes effect.

(f) RUBRICS FOR ALL FRONTIER AI APPLICATIONS.—The Administrator shall consider assigning a scoring factor for each of the criteria listed below for all types of frontier AI applications—

(1) the applicant's plan for securing liability insurance or otherwise mitigating the risks posed by the applicant's AI systems;

(2) the applicant's plan for detecting and reporting incidents and accidents related to its frontier AI systems or hardware, and for learning from such events and adapting so as to minimize the chance that such events will reoccur;

(3) the applicant's demonstrated ability to accurately forecast the capabilities and risks of their frontier AI systems or hardware;

(4) the applicant's resources, abilities, reputation, and willingness to successfully execute the plans described in subsections (f) through (j).

(g) RUBRICS FOR HARDWARE FRONTIER AI APPLICATIONS.—The Administrator shall consider assigning a scoring factor for each of the criteria listed below for hardware frontier AI applications—

(1) the applicant's plan for ensuring that it is aware of the real identities of its customers and that it does not rent or sell hardware to irresponsible or unknown persons;

(2) the applicant's plan for preventing third parties from stealing access to its hardware, e.g., via hacking; and

(3) the applicant's plan for preventing third parties from physically stealing its hardware.

(h) RUBRICS FOR MODEL WEIGHTS FRONTIER AI APPLICATIONS.—The Administrator shall consider assigning a scoring factor for each of the criteria listed below for model weights frontier AI applications—

(1) the applicant's plan for ensuring that it is aware of the real identities of its customers and that it does not rent or sell access to the services made possible by its model weights to irresponsible or unknown persons;

(2) the applicant's plan for preventing third parties from stealing access to its model weights, e.g., via hacking;

(3) the applicant's plan for limiting access to its model weights to persons who also hold a valid model weights permit; and

(4) the extent to which the applicant's intended uses of the model weights pose no major security risks.

(i) RUBRICS FOR TRAINING FRONTIER AI APPLICATIONS.—The Administrator shall consider assigning a scoring factor for each of the criteria listed below for training frontier AI applications—

(1) the extent to which the applicant has clearly specified a maximum intended level of capabilities for the AI system to be trained;

(2) the extent to which the applicant has convincingly explained why the level of capabilities it intends to train will be safe to train;

(3) the extent to which the applicant has developed a theory predicting how the capabilities of its AI system will increase during training as the compute, data, and parameters included in the AI system are scaled up;

(4) the applicant's plan for promptly and reliably detecting all significant discrepancies between the rate at which capabilities increase during training and the rate of increase predicted by the applicant's theory;

(5) the applicant's plan for promptly and fully halting all further training upon discovering a discrepancy between the predicted and actual rate of increase in capabilities;

(6) the applicant's plan as to how, upon obtaining any anomalous results such as an unexpectedly rapid increase in capabilities, the applicant will communicate such anomalies to the Administration, will jointly interpret any anomalous results with the Administration, and will ensure that training does not resume unless and until the applicant and the Administration jointly devise a plan for safely resuming training;

(7) the extent to which the applicant has developed a total compute budget that clearly indicates the maximum amount of compute that will be used for the final training run;

(8) the applicant's plan for ensuring that it will not exceed its total compute budget;

(9) the applicant's plan for ensuring that the AI system it is training will not escape during training or otherwise significantly influence the world outside of the laboratory in which it is being trained;

(10) the extent to which the applicant has developed a timeframe during which it will conduct all training, and the applicant's plan for completing all training within that timeframe;

(11) the extent to which the applicant has outlined the relevant information, data, and programs related to frontier AI that it intends to share during training;

(12) the applicant's plan for which specific persons or job roles will receive each type of information, data, or program that it will share during training; and

(13) the applicant's plan for preventing, detecting, and responding to unauthorized access to its AI systems, including elements of physical security, cybersecurity, and personnel security.

(j) RUBRICS FOR DEPLOYMENT FRONTIER AI APPLICATIONS.—The Administrator shall consider assigning a scoring factor for each of the criteria listed below for deployment frontier AI applications—

(1) the extent to which the applicant provides convincing evidence that the AI system is robustly aligned, i.e., that it will behave as intended across all plausible conditions under which it might be used;

(2) the applicant’s plan for ensuring the traceability of products and services enabled by the applicant’s frontier AI system;

(3) the applicant’s plan for ensuring that its frontier AI system will not be used, accessed, or reverse engineered in countries that lack adequate AI safety legislation;

(4) the applicant’s plan for ensuring that its frontier AI system will not be fine-tuned, connected with plug-ins or utilities, or otherwise modified in such a way as to significantly increase the major security risks posed by that frontier AI system;

(5) the applicant’s plan for monitoring its AI system and for retaining the capability to promptly and fully disable access to that AI system;

(6) the extent to which the applicant’s frontier AI systems would be likely to exacerbate major security risks by accelerating the pace at which new AI capabilities are developed;

(7) the extent to which the applicant’s frontier AI systems could autonomously survive, replicate, or spread; and

(8) the extent to which the applicant’s frontier AI systems directly pose major security risks by contributing to activities such as bioweapons development, nuclear weapons development, automated hacking, autonomous weapons systems, or similar threats.

(k) CONSIDERATIONS FOR OPEN SOURCE FRONTIER AI SYSTEMS.—A scoring factor created under this section shall be designed and applied so as to fairly consider both the risks and benefits associated with open source frontier AI systems, including both the risk that an open source frontier AI system might be difficult or impossible to remove from the market if it is later discovered to be dangerous, and the benefits that voluntary, collaborative, and transparent development of AI offers to society.

(1) OPEN SOURCE CODE.—When the Administrator determines that an open source AI system poses major security risks, the Administrator shall consider whether it is useful and appropriate to impose a partial limitation on access to that open source AI system. For example, the Administrator might impose a partial limitation by—

(A) requiring that the open source project must verify the real identity of anyone wishing to download the source code or model weights or both;

(B) requiring that the open source project confirm that a person accessing the project’s resources has a legitimate, pro-social interest supporting that access, such as contributing to the open source project, conducting research, conducting safety testing, or engaging in entrepreneurship;

(C) requiring that the open source project confirm that a person accessing the project's resources is a responsible actor who can be relied on not to further distribute the source code or model weights or both without permission; or

(D) requiring some combination of the above.

(2) DEGREE OF OPEN SOURCE.—When the Administrator determines that an open source AI system poses major security risks, the Administrator shall consider what degree and kind of restriction of access to the open source AI system will be sufficient to protect the public safety. For example, the Administrator might restrict access to—

(A) only the system's source code;

(B) only the system's model weights;

(C) only the system's training data; or

(D) some combination of the above.

(3) NO AUTOMATIC DETERMINATIONS.—A frontier AI system shall not be considered inherently dangerous or inherently safe based solely on the fact that one or more aspects of the system are open source; instead, the Administrator shall fairly evaluate the system pursuant to the rubrics and procedures in Sections 8 through 11 and shall determine whether the open-source AI system should be permitted based on the extent to which the system poses major security risks.

(4) NO APPLICATION TO NON-FRONTIER OPEN SOURCE.—The rules in paragraphs (1) through (3) above apply only to high-concern and extremely high-concern AI systems. This law does not impose any special requirements on open source AI systems that are classified as low-concern or medium-concern.

SEC. 9. ADDITIONAL REQUIREMENTS FOR EXTREMELY HIGH-CONCERN AI.

(a) DEVELOPMENT OF RISK CRITERIA.—No later than 12 months after the enactment of this Act, the Administrator shall promulgate detailed standards for determining whether an AI system is extremely high-concern.

(b) INTERIM DETERMINATIONS.—Before the Administration has promulgated detailed standards as required by section 9(a), the Administrator may issue a finding that a particular AI system that would otherwise qualify as high-concern is instead extremely high-concern. Any such finding shall be published in the Federal Register, along with an explanation of which major security risks are posed by that AI system.

(c) DEVELOPMENT OF SAFETY CRITERIA.—No later than 30 months after the enactment of this Act, the Administrator shall promulgate detailed standards for determining whether an extremely high-concern system is nonetheless safe enough to be trained, and for determining whether an extremely high-concern system is nonetheless safe enough to be deployed. Such standards should include a set of assessments that substantially cover all of the following topics—

(1) whether the AI system's architecture is fundamentally safe;

(2) whether there are mathematical proofs that the AI system is robustly aligned;

(3) whether the AI system has been specifically found to be inherently unable to assist with the development of biological, chemical, nuclear, and radiological weapons;

(4) whether the AI system has been specifically found to be inherently unable to autonomously replicate or spread; and

(5) whether the AI system has been specifically found to be inherently unable to accelerate scientific or engineering progress to such a degree as to pose national security risks or destabilize the global balance of power.

(d) EVALUATION OF SAFETY FOR EXTREMELY HIGH-CONCERN AI SYSTEMS.

(1) Before the promulgation of the standards in section 9(c), it is unlawful to train or deploy any extremely high-concern AI system.

(2) After the promulgation of the standards in section 9(d), the Administrator shall develop additional application forms, rubrics, and processes to evaluate whether a proposed training run or deployment is safe enough to train or deploy despite the extremely high risks involved.

(3) In evaluating each such training run or deployment, the Administrator shall not determine that an extremely high-concern AI system is safe solely based on a lack of evidence that the AI system poses major security risks. Instead, any such determination may be made only after the Administrator finds, based on clear and convincing evidence, that one or more features of the AI system affirmatively rules out any significant possibility that the AI system could pose major security risks.

SEC. 10. ADJUDICATION OF PERMIT APPLICATIONS.

(a) **APPOINTMENT OF AILJs.**—The Administrator shall appoint AI Judges (AIJs) who shall have competent scientific ability and sufficient legal knowledge to faithfully and accurately apply the laws and regulations under this Act. AIJs shall be entitled to compensation as if they were administrative law judges under pay scale AL-3.

(b) **INITIAL REFERRAL.**—Upon receipt of a frontier AI permit application, the Administrator shall refer the application to a panel of three randomly selected AIJs and shall forward a copy of the application to the Deputy Administrator for Public Interest.

(c) **USE OF PRIVATE TECHNICAL EVALUATORS.**—The Administrator is encouraged to make use of private technical evaluators, i.e., persons who do not work directly for the government and who will provide an accurate technical description of the capabilities, features, risks, and security mechanisms associated with a particular AI system. Where feasible, private technical evaluators should be used in preference to government personnel in order to provide this type of description and analysis. Such evaluators may be hired by the Administrator to assist with the adjudication of applications, and such evaluators may be hired by applicants to assist with the submission of applications. A private technical evaluator must have all of the following qualifications, as well as any additional qualifications established by the Administrator by regulation:

(1) they are legally independent from the owners and developers of the systems they are analyzing, with no significant conflict of interest or overlap in their management and board of directors.

(2) any payment or equity the evaluator receives from an owner or developer for the evaluation of an AI system will make up less than 25% of the evaluator's annual revenue.

(3) the evaluator does not have any financial interest in the owner or developer of the AI system beyond its fee for evaluating the AI system.

(4) the amount of the evaluator's fee for evaluating the AI system is not tied to the content of the evaluator's findings.

(5) the evaluator has the requisite technical skill to competently evaluate the features and risks of the AI system.

(6) the evaluator does not make any final recommendations as to the suitability or safety of the AI system, but instead confines themselves to describing the features of the system so that a final recommendation can be made by the Administration.

(d) RECOMMENDATION BY AIJ.—Within 60 days after receipt of an application, a panel of AIJs shall, by majority vote, take exactly one of the following actions with respect to that application—

(1) recommend unconditional approval of the application;

(2) recommend approval of the application with conditions;

(3) recommend that specific revisions be made to the application and that the application then be resubmitted for reconsideration;

(4) recommend rejection of the application; or

(5) return the application without a recommendation due to the AIJ's inability to process the volume of applications that the AIJ has received.

(e) RESOLUTION OF CONFLICTING RECOMMENDATIONS.—If there is no majority supporting the same option in section 10(d) above, then the AIJs shall briefly meet and confer on their recommendations. If they still cannot agree within 5 days after their initial vote, then the application is considered rejected.

(f) BASIS FOR DECISION BY AIJ.—In making each recommendation, an AIJ shall apply the rubrics developed under sections 8 and 9 to each application, and shall consider the extent to which granting a permit would increase major security risks, and shall weigh such risks against the utility or value provided by the frontier AI activities subject to the permit. The AIJ shall recommend rejection of any application which either fails the applicable rubric. The AIJ shall also recommend rejection of any application which, in the AIJ's opinion, would significantly increase major security risks during the time period covered by the permit.

(g) OPINION BY AIJ.—An AIJ who formed part of the majority for each recommendation shall provide a short written opinion explaining the basis for the AIJ's recommendation and evaluating the estimated effect of the application's approval on major security risks. The opinion shall be provided to the Administrator, the Deputy Administrator for Public Interest, and the applicant within 5 days after the recommendation is made.

(h) FINALITY OF RECOMMENDATIONS.—A recommendation becomes final 20 days after it is provided to the Administrator, the Deputy Administrator for Public Interest, and the applicant if none of those three persons have appealed the recommendation within that time

period. If the recommendation included conditions, then the recommendation is only deemed finally approved if the applicant accepts those conditions. The Administrator shall promptly issue or renew a permit application that has received a final recommendation of approval. The applicant is entitled to have such a permit issued or renewed.

SEC. 11. APPEALS OF PERMIT APPLICATIONS.

(a) **TIMING AND PARTIES FOR APPEAL.**—The applicant, the Deputy Administrator for Public Interest, or the Administrator may appeal an AIJ’s recommendation to the AI Appeals Board by filing a notice of appeal within 20 days of receiving the AIJ’s written opinion. The notice of appeal shall consist of a short, plain statement of the facts and reasoning supporting the appeal. An applicant who has not received a recommendation or return within 70 days after submitting an application may likewise file an appeal to the AI Appeals Board protesting the delay.

(b) **COMPOSITION OF APPEALS BOARD.**— The Artificial Intelligence Appeals Board shall consist of seven Board Members selected by the Administrator from a list of qualified candidates prepared by the Deputy Administrator for Public Interest.

(1) **QUALIFICATIONS.**—Each Board Member shall be a highly qualified professional with relevant expertise and no record of disciplinary sanction in their field(s). The Administrator should assemble a Board with a diverse set of professional strengths. A Board Member who qualifies based on legal expertise shall be a member in good standing of a State Bar or the Bar of the District of Columbia and shall have demonstrated interest and proficiency in the application of law to AI. A Board Member who qualifies based on scientific expertise shall have published original research in the fields of computer science, artificial intelligence, or electrical engineering. A Board Member who qualifies based on risk management or national security expertise shall have appropriate certifications for their field and shall have experience in dealing with AI-specific risks or experience with a wide range of risks, including financial, operational, strategic, and compliance-related risks.

(2) **RECUSAL.**—It is the responsibility of each Board Member to recuse themselves when that Board Member has a real or apparent conflict of interest for an appeal. In addition, the Deputy Administrator for Public Interest may petition the Administrator to recuse a member of the Appeals Board on the basis of a real or apparent conflict of interest. The recusal of one or more Board Members does not prevent the Board from achieving quorum.

(3) **BOARD COUNSEL.**—An Appeals Board that lacks sufficient expertise in any field relevant to a question before it may appoint and consult legal counsel or expert advisors.

(c) **PROCEDURE FOR APPEALS BOARD.**—The Appeals Board shall consider *de novo* all questions of law presented by each appeal. The Appeals Board may likewise consider *de novo* any factual question presented by an appeal, or the Appeals Board may apply an abuse of discretion standard to one or more factual questions resolved by the AIJs, at the discretion of the Appeals Board. The Appeals Board shall resolve each appeal before it by majority vote of the participating Board Members. In case of a tie, an application is considered rejected. The Deputy Administrator for Public Interest may attend all meetings of the Appeals Board, may present arguments, and may ask questions, but shall not vote or preside over the

meetings. The Appeals Board shall randomly select one of its members to preside over the appeal; if that member voted with the majority, then that member shall write an opinion summarizing the result and reasoning of the appeal, and otherwise that member shall delegate the writing of such opinion to a Board Member who voted with the majority. The Appeals Board shall resolve each appeal and provide a copy of its opinion to the Administrator, the Deputy Administrator for Public Interest, and the applicant within 60 days of receiving the applicant's notice of appeal.

(d) INTERVENTION BY ADMINISTRATOR.—If the Administrator disagrees with the opinion of the Appeals Board, the Administrator may modify or reverse that opinion within 10 days of the Administrator's receipt of the opinion by providing a written explanation of that disagreement that explains, in detail, why and how the Appeals Board's decision fails to adequately further the purposes of this Act.

(e) EXHAUSTION OF ADMINISTRATIVE REMEDIES.—All administrative remedies with respect to a licensing application are deemed to be exhausted—

(1) when the Administrator issues a statement under paragraph (3) of this subsection, or

(2) 10 days after a copy of the Appeals Panel's opinion is provided to the Administrator, if the Administrator has not yet issued a statement under section 11(d).

(f) JUDICIAL REVIEW.—After all administrative remedies have been exhausted, either the applicant or the Deputy Administrator for Public Interest or both may appeal a licensing application to the Court of Appeals for the District of Columbia Circuit. The Deputy Administrator for Public Interest shall have standing for such an appeal as a representative of the public's interest in mitigating major security risks. No such appeal may be filed more than 20 days after the exhaustion of administrative remedies.

(g) STATUS OF PERMITS DURING REVIEW.—No applicant for a frontier AI permit may take any action for which such a permit is required while that permit is pending adjudication, administrative appeal, or judicial appeal, unless the applicant is applying for a renewal permit, and the action was permitted by the terms of the applicant's most recent permit.

SEC. 12. HARDWARE MONITORING.

(a) SELF-REPORTING FOR HIGH-PERFORMANCE AI HARDWARE.—The Administrator shall monitor the status of high-performance AI hardware based on a self-reporting requirement, as follows:

(1) No later than 90 days after this law is enacted, the Administrator shall create a website containing a form to be filled out by any person who buys, sells, gifts, receives, trades, destroys, or transports one or more units of high-performance AI hardware. (For this section, a person "transports" such a unit if they cause its location to change by at least 10 miles.)

(2) The Administrator shall advertise the availability of the form by publishing notice in the Federal Register, by prominently displaying the form's availability on the Administration's website, and through at least one other method that the Administrator finds proper and useful to alert the public that the form is available.

(3) Beginning 60 days after the publication described in subsection (a)(2), all persons who buy, sell, gift, receive, trade, or transport one or more high-performance AI hardware units must record that transaction using the form.

(4) A person who acquires or transports a high-performance AI hardware unit must complete the form by the earlier of—

(A) 24 hours after the person first uses one or more of the newly acquired or transported high-performance AI hardware units; or

(B) 10 days after the person acquires or transports the high-performance AI hardware units.

(5) A person who destroys or transfers a high-performance AI hardware unit must complete the form within 10 days after destroying or transferring the unit.

(6) A completed form must include information about the identity of the previous owner, the identity of the new owner, the type of high-performance AI hardware, the quantity of high-performance AI chips, and the total estimated throughput of the high-performance hardware being transferred.

SEC. 13. ANALYSIS AND REPORTING.

(a) TABULATION OF HARDWARE REPORTS.

(1) No later than the 10th day of each month, the Administrator shall compile the data acquired via reports pursuant to Section 12 and attempt to identify and describe all of the following:

(A) The distribution of high-performance AI hardware by geography, industry, and type of owner.

(B) The most notable concentrations of high-performance AI hardware, especially collections of high-performance AI hardware in the hands of persons who do not have a current frontier AI hardware permit.

(C) Patterns in the flow and stockpiles of high-performance AI hardware.

(D) Notable changes in the flow of high-performance AI hardware over time.

(2) No later than the 15th day of each month, the Administrator shall collect and compile information on all of the following:

(A) The amount of compute authorized to be possessed by each holder of a hardware frontier AI permit.

(B) The total amount of compute provided by all high-performance AI hardware.

(C) The rate at which high-performance AI hardware is being manufactured, expressed in terms of the amount of compute being created by such manufacturing.

(D) The rate at which the effective power of the total supply of compute is increasing due to improvements in algorithmic efficiency.

(E) The primary purposes for which high-performance AI hardware is being used.

(3) No later than the 20th day of each month, the Administrator shall collate and analyze the information collected via paragraphs (7) and (8) with the goal of determining which high-performance AI hardware (if any) is not adequately accounted for.

(b) PROACTIVE ANALYSIS OF THREATS.—The Administrator shall proactively attempt to detect, identify, and understand the most important sources of major security risks from frontier AI systems. The Administrator shall use the powers under this Act to reduce and mitigate those risks. If the Administrator detects a major security risk from frontier AI systems that the Administration lacks the power to adequately mitigate, then the Administrator shall immediately so inform the National Security Advisor, the Director of the Cybersecurity & Infrastructure Security Agency, the Director of the Centers for Disease Control and Prevention, and the Director of the Federal Emergency Management Agency.

(c) INVESTIGATIONS.—The Administrator may, in its discretion, make such investigations as it deems useful to fulfill any of the Administrator’s obligations under this Act, including investigations—

(1) to support the proactive analysis of threats described in subsection (b), or

(2) to determine whether any person or entity has violated, is violating, or is about to violate any provisions of this Act, the rules or regulations, or a permit issued thereunder.

(d) TAKING OF EVIDENCE.—For the purpose of any such investigation, or any other proceeding under this Act, the Administrator or any officer designated by the Administrator is empowered to administer oaths and affirmations, subpoena witnesses, compel their attendance, take evidence, enter onto premises where frontier AI systems or frontier AI hardware are believed to exist, and require productions of any books, papers, correspondence, memoranda, or other records the Administrator deems relevant or material to the inquiry. Such attendance of witnesses and the production of any such records may be required from any place in the United States or any State at any designated place of hearing. The Administration shall pay the reasonable expenses of such attendance and production. The refusal of any person to fully cooperate with such taking of evidence may be punished according to Federal law, including civil and criminal penalties for contempt of court.

(e) ANNUAL BULLETIN IN FEDERAL REGISTER.—No later than April 1st of each year, the Administrator shall publish a bulletin in the Federal Register, which shall be current as of March 15th of that year, and which shall include all relevant information in each of the following categories that has not previously been published in such a bulletin:

(1) The number of persons employed pursuant to section 4(c)(3) of this Act.

(2) Deputy Administrators who have been removed for cause, and a description of each such cause.

(3) Waivers of conflict of interest that have been granted, together with a description of the reasons for such waivers.

(4) Recommendations as to frontier AI licensing that the Administrator has modified or reversed, together with an explanation of why the recommendations were modified or reversed.

(5) Each use of emergency powers under Section 11, together with an explanation of why the use was thought necessary and the current status of the event or activity that was subject to the use of emergency powers.

(f) ANNUAL REPORT TO CONGRESS.—No later than October 1st of each year, the Administrator shall submit a report to Congress.

(1) SIZE.—The report shall contain no more than 20 letter-sized pages, using a reasonable font size and typesetting, including all attachments, prefaces, and exhibits.

(2) CONTENTS.—The report shall inform Congress about the most important major security risks posed by frontier AI systems, what the Administration is doing to address those risks, which of those risks (if any) the Administration lacks the power to adequately address, and what actions (if any) the Administrator recommends that Congress take in order to reduce those risks to an acceptable level.

(g) QUARTERLY BRIEFINGS FOR OSTP.—No later than the 30th day of March, June, September, and December of each year, the Administrator or a Deputy Administrator shall meet in person with an official at the White House Office of Science and Technology Policy and brief that official on major security risks from AI.

(h) OTHER REPORTS TO GOVERNMENT.—From time to time, the Administrator shall make other reports to Congress, to the White House, and to executive agencies that inform and educate them about major security risks from AI, especially when—

(1) such risks have recently increased,

(2) such risks are likely to increase soon, or

(3) a government official is about to take an action or make a decision related to such risks.

SEC. 14. CIVIL LIABILITY.

(a) WHO OWES DUTY OF CARE.—All persons engaged in the development or use of artificial intelligence owe a duty of care to exercise appropriate caution. The duty of care is owed to all persons who are residents of the United States. This duty of care is owed by any person who knowingly owns, possesses, trains, develops, deploys, or uses any of the following—

(1) artificial intelligence,

(2) any specialized hardware that is designed to support AI, or

(3) the model weights for any frontier AI.

(b) DUTY OF CORPORATE PARENT OR SENIOR CORPORATE OFFICER.—A person owes the duty of care described in subsection (a) if that person has both—

(1) the authority to control the behavior of a person who owes the duty of care described in subsection (a), and

(2) actual or constructive knowledge that the person who owes the duty of care described in subsection (a) is violating that duty of care or is likely to violate that duty of care.

(c) OBLIGATIONS UNDER DUTY OF CARE.—Persons who are subject to this duty of care have an affirmative obligation to ensure each of the following:

(1) None of their AI systems cause harm to innocent bystanders, i.e., to persons who are not customers, users, or developers of the AI and who have not maliciously interfered with the AI.

(2) Their frontier AI systems do not escape or spread to third party hardware whose owners have not affirmatively consented to host that frontier AI.

(3) The model weights of their frontier AI systems are not leaked, stolen, or otherwise made publicly available.

(4) Their frontier AI systems are reasonably secure against misuse by third parties, which includes the obligation to—

(A) attempt to identify the most important avenues for misuse of their frontier AI,

(B) monitor their frontier AI for potential misuse, and

(C) upon becoming aware of a third party's misuse or credible threat to misuse their frontier AI, immediately deny that third party access to their frontier AI.

(5) Their high-performance AI hardware is reasonably secure against misuse by third parties, which includes the obligation to—

(A) attempt to identify the most important avenues for misuse of their high-performance AI hardware,

(B) monitor their high-performance AI hardware for potential misuse, and

(C) upon becoming aware of a third party's misuse or credible threat to misuse their high-performance AI hardware, immediately deny that third party access to their high-performance AI hardware.

(d) JOINT AND SEVERAL LIABILITY.—All persons who have violated the duty of care imposed by this Act with respect to the same AI system, high-performance AI hardware, or set of model weights are jointly and severally liable for any violations of that duty of care that contributed to the same harm or to a set of substantially related harms.

(e) PRIVATE RIGHT OF ACTION.—A person, group of people, or putative class who allege specific facts that plausibly suggest a claim for at least \$100 million in tangible damages based on a violation of the duty of care established by this section or based on the strict liability established under this section shall have a private right of action and may bring suit for those damages, together with costs of suit and reasonable attorneys' fees, in any federal district court that has personal jurisdiction and venue under Title 28 of the United States Code. For claims of less than \$100 million in tangible damages, this section is not intended to create, destroy, or modify any private rights of action.

(1) QUALIFYING DAMAGES.—Damages are considered “tangible” if they are wrongful death, physical injury or illness, direct financial losses, conversion, the lost value of destroyed or corrupted data, payments made in response to ransomware attacks, or damage to physical property or real estate.

(2) EXCLUDED DAMAGES.—Damages are not considered “tangible” if they are emotional distress, libel, slander, invasion of privacy, consequential damages, loss of goodwill, loss of business opportunities, or violations of intellectual property rights.

(f) PUBLIC RIGHT OF ACTION.—The Administrator may sue any person who violates the duty of care established by this section or who is strictly liable under this section.

(1) CONTENT OF LAWSUIT.—Such a lawsuit may include any of the following—

(A) a request for injunctive or equitable relief,

(B) an attempt to recover damages on behalf of identified victims for distribution to those victims; and

(C) an attempt to recover damages on behalf of the public at large.

(2) CIVIL PENALTY.—If the Administrator is successful in such a lawsuit, the Court shall assess a civil penalty of at least \$25,000 and no more than \$500,000 per defendant, payable to the Treasury, taking into account the degree to which each defendant has contributed to major security risks and the profit, if any, that each defendant derived from the violation. A defendant who knowingly continued to violate the duty of care may be assessed an additional civil penalty of up to \$100,000 per day of the continuing violation.

(3) PROSECUTION OF LAWSUIT.—The Administrator may directly prosecute such a lawsuit, or may refer such a lawsuit to the Department of Justice for prosecution.

(g) NEGLIGENCE *PER SE*.—In any lawsuit alleging a violation of the duty of care created by this subsection, a defendant shall be deemed negligent *per se* if any of the following apply:

(1) The defendant was required to obtain a frontier AI permit and failed to do so.

(2) The defendant violated the terms of their frontier AI permit.

(3) The defendant made a material misrepresentation (including a misrepresentation by omission) on their application for a frontier AI permit.

(h) STRICT LIABILITY.—In any civil lawsuit where the plaintiff or plaintiffs have alleged specific facts that plausibly suggest a claim for at least \$100 million in tangible damages caused by a frontier AI system, each defendant shall be strictly liable for all tangible damages caused by any event that arises out of or meaningfully relates to that frontier AI system.

(1) SUBSTANTIAL FACTOR CAUSATION.—If strict liability applies, then a plaintiff is not required to prove that a defendant's actions were the proximate cause of the plaintiff's harm. Instead, a plaintiff may prove that a defendant's actions were a substantial factor in causing the plaintiff's harm.

(2) NO REQUIREMENT TO SHOW DEFECT.—If strict liability applies, then a plaintiff is not required to prove that any aspect of a defendant's AI was defective.

(i) EXCEPTIONS FOR BONA FIDE ERROR.—The provisions of subsections (g) and (h) shall not apply to a defendant who shows by a preponderance of the evidence that any violation of the duty of care established by this section was unintentional and resulted from a bona fide error notwithstanding the maintenance of procedures reasonably adapted to avoid any such error. Bona fide errors include errors that are solely due to clerical errors, arithmetic errors,

or printing errors. An error of legal judgment or technical judgment with respect to a person's obligations under this statute is not a bona fide error.

(j) NO DEFENSE BASED ON OPEN SOURCE.—It shall not be a defense or excuse for any civil liability under this section that a defendant's role was limited to hosting, developing, fine-tuning, or distributing a product or service that is free, collaborative, or open source.

(k) FORESEEABILITY OF MISALIGNMENT.—It shall not be a defense or excuse for any civil liability under this section that a defendant was unable to foresee the precise manner in which a frontier AI system would become misaligned or unreliable. As a matter of law, persons developing frontier AI are deemed to have foreseen the general possibility that an apparently well-aligned AI system may turn out to exhibit undesirable behavior after being scaled up, more widely deployed, fine-tuned, connected to additional utilities and plug-ins, or otherwise placed into a more dangerous environment. Persons who train or deploy frontier AI are liable for such misbehavior.

(l) PUNITIVE DAMAGES.

(1) WHEN AVAILABLE.—Punitive damages shall be awarded whenever a defendant is held liable for a violation of the duty of care established by this section, if that defendant—

(A) recklessly engaged in misconduct while knowing that this misconduct had the potential to cause major security risks; or

(B) recklessly engaged in misconduct that narrowly avoided causing major security risks.

(2) AMOUNT.—In setting the amount of such punitive damages, a court shall take into account the high value to society of avoiding major security risks, and shall award an amount of punitive damages that is sufficient to deter future violations. In the absence of evidence to the contrary, an award of nine times the value of the compensatory damages shall be considered to balance society's interest in avoiding major security risks with the requirements of due process.

(3) REPREHENSIBILITY.—A plaintiff is not required to prove malice or oppression in order to receive punitive damages based on this subsection. Instead, a plaintiff may demonstrate that a defendant's tolerance for major security risks showed complete indifference to the safety of the public and is therefore sufficiently reprehensible to justify punitive damages.

SEC. 15. CRIMINAL LIABILITY.

(a) PAYMENT OF FINES.—Any person who commits any crime in this subsection shall be fined according to the schedule in 18 U.S. Code § 3571 based on the classification of the crime.

(b) FELONY PERMIT VIOLATIONS.—Any person who performs any of the following activities shall be guilty of a class C felony, and, upon conviction, shall be imprisoned for less than 25 years but more than 10 years:

(1) The person has received an emergency order under Section 11 to cease an activity related to frontier AI, and the person fails to take all steps within the person's power to promptly and fully comply with that order.

(2) The person's application for a frontier AI permit has been rejected, and the person nevertheless conducts activities of the type that were contemplated by that application.

(3) The person's application for a frontier AI permit was approved with conditions, and the person conducts activities of the type contemplated by the application while knowingly violating those conditions.

(4) The person makes, approves, or submits any material statement of fact on an application for a frontier AI permit while having actual knowledge that the statement is false.

(5) The person states an intention to take a safety precaution or otherwise mitigate a risk on an application for a frontier AI permit that is fraudulent in that the person did not intend to take that safety precaution or otherwise mitigate that risk at the time the statement was made.

(c) MISDEMEANOR PERMIT VIOLATIONS.—Any person who performs any of the following activities shall be guilty of a class A misdemeanor, and, upon conviction, shall be imprisoned for less than 1 year but more than 6 months:

(1) The person knowingly improves or uses the capabilities of frontier AI, and such activity requires a frontier AI permit, and the person does not have a currently valid frontier AI permit that covers this improvement or use.

(2) The person is obligated to take or refrain from an action under the terms of a frontier AI permit, and the person recklessly fails to take or refrain from that action.

(3) The person makes, approves, or submits any part of an application for a frontier AI permit while having actual or constructive knowledge that such part of the application is significantly incomplete or misleading.

(4) The person knowingly alters or adjusts an AI system so as to artificially reduce the AI system's performance on a benchmark or test without similarly reducing the AI system's true capabilities, thereby causing the AI system to receive less regulatory scrutiny.

(d) SELF-REPORTING INFRACTIONS.—Any person who is required to self-report a transaction involving high-performing AI hardware and who fails to do so within the allotted time shall be guilty of an infraction, and, upon conviction, shall be imprisoned for up to 5 days.

(1) MINIMUM FINE.—The fine for this infraction shall be no less than \$5,000 or twice the total price of the high-performance AI hardware, whichever is greater.

(2) REPEAT OR SERIOUS VIOLATIONS.—However, if the person was previously convicted under this paragraph prior to the date on which the new infraction was committed, or if the first infraction involves a failure to report at least 20 times the compute required to trigger the reporting requirement, then the person shall instead be guilty of a class A misdemeanor, and, upon conviction, shall be imprisoned for less than 1 year but more than 6 months.

(e) OTHER CRIMES.—Any person who recklessly violates any other provision of this Act, or any rule or regulation thereunder, shall be guilty of a class B misdemeanor, and, upon conviction, shall be imprisoned for less than 6 months but more than 30 days.

(1) WILLFUL FAILURE OF OFFICER.—A person working for the Administration who fails to complete a duty prescribed by the Act shall not be guilty under this subsection unless that person had the resources to perform the duty and willfully and intentionally refused to perform it.

(2) ELEVATION TO FELONY.—A person shall instead be guilty of a class D felony, and shall be imprisoned for less than 10 years but more than 5 years, if at least two of the following are true:

(A) In committing the violation, the person acted or failed to act intentionally or with actual knowledge.

(B) The person's violation was likely to significantly increase major security risks, or did significantly increase major security risks, or caused at least \$100 million in tangible damages.

(C) The person was convicted of any misdemeanor or felony under this Act prior to the date on which the new violation was committed.

(f) CRIMINAL ATTEMPT.—A person who attempts to commit a federal offense as defined in this section shall be subject to a penalty one degree lower than that prescribed for the completed offense. For example, an attempt to commit an offense classified as a Class B misdemeanor under this statute shall incur penalties as specified for a Class C misdemeanor.

(g) NO CORPORATE INDEMNIFICATION.—Whenever a fine under this Act is imposed upon any officer, director, employee, agent or stockholder of an entity, such fine may not be paid, directly or indirectly, by such entity. It is unlawful for an entity to increase the compensation paid to such an agent in such a way as to relieve the burden of a fine imposed under this Act.

(h) FINES NOT LIMITED BY PROFITS.—When any entity is fined under this subsection, and the entity is a non-profit corporation or otherwise cannot be adequately deterred with a fine based on the amount of the entity's profits, then the Court may increase the fine imposed on each such entity up to \$2 million for a misdemeanor, or up to \$25 million for a felony, taking into account the seriousness of the offense and the need for adequate deterrence. This paragraph shall not be construed to limit the maximum amount of a fine imposed under any other provision of Federal law.

(i) ADDITIONAL CORPORATE PENALTIES.

(1) FOR MISDEMEANORS.—When a Court sentences an entity based on a misdemeanor under this Act, the Court shall suspend all frontier AI permits held by that entity for a period of less than 1 year but more than 1 month. An entity with a suspended frontier AI permit shall prevent its users and customers from accessing its frontier AI systems during the period of such suspension, and shall not conduct any research, development, or testing of its frontier AI systems during the period of such suspension.

(2) FOR FELONIES.—When a Court sentences an entity based on a felony under this Act, the Court shall cancel all frontier AI permits held by that entity, and the entity shall be ineligible to apply for frontier AI permits for a period of 5 years, and the Court shall

order the entity to immediately encrypt all frontier AI model weights in the entity's possession with a key held by the Administration, which key shall be used to decrypt the model weights if and only if the model weights are sold or transferred to a person with a valid frontier AI permit for said model weights. The entity shall sell or transfer all of its frontier AI hardware to other person(s) with a valid frontier AI permit within 60 days of the sentence; any frontier AI hardware that cannot be sold within that time must be destroyed.

(j) PROSECUTION OF CRIMES.—The Administrator may prosecute any crime under this section directly, or the Administrator may refer any crime under this section to the Department of Justice for prosecution.

(k) STATUTE OF LIMITATIONS.—The statute of limitations for all felonies under this Act is 10 years. The statute of limitations for all misdemeanors under this Act is 5 years. The statute of limitations for all infractions under this Act is 2 years.

SEC. 16. EMERGENCY POWERS.

(a) WHEN EMERGENCY POWERS ARE AVAILABLE.—The emergency powers provided by this section shall be available whenever—

(1) the President by proclamation or Executive order declares a national emergency to exist by reason of a major security risk related to frontier AI; or

(2) the Administrator determines that one or more frontier AI systems pose a clear and imminent major security risk that cannot be reliably prevented through ordinary civil and criminal enforcement, and publishes that determination in a formal declaration.

(b) DURATION OF EMERGENCY POWERS.—If the Administrator initiates the use of emergency powers under this section, the emergency powers shall remain in effect for no more than 60 days unless they are confirmed by the President or the Congress of the United States. If the President initiates or confirms the use of emergency powers under this section, the emergency powers shall remain in effect for no more than 1 year unless they are confirmed by Congress.

(c) SCOPE OF ADMINISTRATOR'S EMERGENCY POWERS.—The Administrator may take any or all of the following actions pursuant to emergency powers under this section—

(1) immediately suspend a frontier AI permit;

(2) issue a cease-and-desist order instructing a person not to take an action related to frontier AI;

(3) issue an order instructing a person to take a safety precaution related to frontier AI;

(4) seize, sequester, or encrypt model weights used or designed or intended for use in frontier AI systems;

(5) issue a restraining order that prevents specified persons from using, accessing, or physically approaching specified frontier AI systems or hardware;

(6) issue a general moratorium on the use or development of frontier AI; and

(7) take any other actions consistent with this statutory scheme that the Administrator deems necessary to protect against an imminent major security risk.

(d) SCOPE OF PRESIDENTIAL EMERGENCY POWERS.—If the President initiates or confirms the state of emergency under this section, then, in addition to the powers listed in subsection (c), the Administrator may also take any or all of the following actions pursuant to emergency powers under this section—

- (1) cancel a frontier AI permit;
- (2) seize or destroy hardware that is used or designed or intended for use in frontier AI systems;
- (3) delete model weights used or designed or intended for use in frontier AI systems; and
- (4) enforce any or all of the above measures by conducting inspections, placing guards, physically removing any unauthorized persons from specified facilities related to AI, or, if necessary to protect against an imminent major security risk, taking full possession and control of specified locations or equipment related to AI.

(e) ANNOUNCEMENT OF EMERGENCY POWERS.—The Administrator shall notify all persons who are directly affected by the use of the Administrator’s emergency powers via personal service or overnight delivery. However, if the Administrator issues a general moratorium under paragraph (c)(6) above, then the Administrator shall instead announce the moratorium via television, radio, and the front page of its website.

(f) TIMING OF EMERGENCY POWERS.—A person is obligated to begin complying with an emergency order under this section from the time that person receives actual notice of the order, or 72 hours after the person is served with notice of the order, whichever comes first.

(g) ENFORCEMENT OF EMERGENCY POWERS.—In order to enforce compliance with emergency orders under this section, the Administrator may request a temporary loan of appropriate personnel from the Federal Bureau of Investigation or the Federal Marshals or both; these agencies shall make reasonable efforts to provide such personnel upon request. In order to enforce compliance with emergency orders under this section, the Administrator is authorized to coordinate, oversee, and direct any or all of the following—

- (1) special agents assigned directly to the Administration;
- (2) personnel loaned by the Federal Bureau of Investigation;
- (3) personnel loaned by the Federal Marshals; and
- (4) any other federal law enforcement officers who are willing to assist.

(h) REVIEW OF EMERGENCY POWERS.—A person who objects to an emergency order issued under this subsection on technical or policy grounds may appeal the order to the Artificial Intelligence Appeals Board, which shall process the appeal following all the procedures specified in Section 11 of this Act. A person who objects to an emergency order as unlawful or unconstitutional may appeal the order to the federal district court having jurisdiction and venue over the matter, as provided by the applicable provisions of Title 28, United States Code.

(i) **STANDARD FOR REVIEW.**—In reviewing an emergency order under this Act, neither the Appeals Panel nor any Court may weaken or set aside that order unless there is clear and convincing evidence of at least one of the following—

- (1) doing so will not pose major security risks, or
- (2) the order was made without adequate legal authority.

(j) **COMPENSATION FOR LOSSES.**—A person who suffers economic losses based on their compliance with an emergency order is entitled to compensation from the United States.

(1) **HOW CALCULATED.**—Such losses shall be calculated based on expenses actually incurred, investments made and lost, and the value of property that has been destroyed. Such losses shall not be calculated based on lost profits, lost goodwill, lost business opportunities, or consequential damages.

(2) **LUMP SUM COMPENSATION.**—With respect to a specific emergency order, Congress may appropriate a sum of money to satisfy all losses incurred based on compliance with that order and direct the Administrator to distribute that money among all who have suffered such losses in proportion to those losses. If Congress does so, all further entitlement to compensation from the United States based on compliance with that emergency order is extinguished.

(3) **REQUIREMENT OF INNOCENCE.**—A person who materially contributed to the need for an emergency order through that person’s negligence or violation of this Act is not entitled to any compensation under this subsection.

SEC. 17. WHISTLEBLOWER PROTECTION.

(a) **WHO QUALIFIES AS FRONTIER AI WHISTLEBLOWER.**—For the purposes of this section, a “frontier AI whistleblower” is defined as any person who has—

- (1) opposed or refused to engage in any practice forbidden by this Act;
- (2) testified, assisted, reported, made allegations in, or otherwise participated in any investigation, litigation, hearing, or proceeding directly related to frontier AI;
- (3) reported the existence or likelihood of any major security risk from frontier AI to an appropriate superior within the whistleblower’s organization; or
- (4) reported the existence or likelihood of any major security risk from frontier AI to the government or to the press after trying and failing to mitigate the risk through internal reporting.

(b) **ACCURACY OF STATEMENTS.**—To qualify for the protections of this section, the statements made by a frontier AI whistleblower must, at the time that they were made—

- (1) have been substantially correct, or
- (2) have been supported by the whistleblower’s reasonable and good faith belief that the statements were substantially correct.

(c) **UNLAWFUL PUNISHMENT.**—It shall be unlawful for an employer to discharge, demote, suspend, threaten, harass, fine, blacklist, discriminate against, or penalize a frontier AI whistleblower in any other way, subject to the following exceptions:

(1) An employer may suspend a whistleblower with full pay for up to one month in order to conduct an investigation.

(2) An employer may penalize a whistleblower based on unrelated conduct, so long as the employer can document with substantial evidence that (A) the whistleblower actually engaged in this conduct, and (B) the penalty imposed was a typical and reasonable response to that conduct.

(3) An employer may lawfully discharge a frontier AI whistleblower by (A) paying that whistleblower two years' salary and benefits in addition to any severance or other awards to which that whistleblower would otherwise be entitled, and (B) providing a neutral reference to other employers who inquire about that whistleblower that is limited to confirming the whistleblower's job title(s) and dates of employment.

(d) PRIVATE REMEDIES.—A frontier AI whistleblower who alleges that they have been unlawfully punished is entitled to pursue all the remedies and procedural advantages of—

(1) 5 USC § 1204 if they are a federal employee, and otherwise;

(2) 18 USC § 1514A.

SEC. 18. INTER-AGENCY COOPERATION.

(a) EXPERT SUPPORT.—Upon request from any other Federal agency for expertise, technical assistance, or other support from the Administration, the Administration may provide that support.

(b) REQUIRED CONSULTATION BY OTHER FEDERAL AGENCIES.—Any Federal agency, including but not limited to the Federal Trade Commission and the Antitrust Division of the Department of Justice, engaged in investigation, regulation, or oversight with respect to the impact of frontier AI Systems on consumer protection, competition, civic engagement, or democratic values and institutions shall consult with the Administration in carrying out that investigation, regulation, or oversight.

(c) REQUIRED CONSULTATION WITH OTHER FEDERAL AGENCIES.—The Administration, in carrying out investigation, regulation, or oversight with respect to the impact of frontier AI Systems on consumer protection, competition, civic engagement, or democratic values and institutions, shall consult with any other Federal agency, including the Federal Trade Commission and the Antitrust Division of the Department of Justice, that is engaged in investigation, regulation, or oversight with respect to the impact of frontier AI Systems on consumer protection, competition, civic engagement, or democratic values and institutions.

(d) AMENDMENT TO ANTITRUST LAWS.—Section 7A of the Clayton Act (15 U.S.C. 18a) is amended by adding at the end the following—

“(l) Frontier AI labs

“(1) In this subsection—

“(A) the terms ‘Administration’ and ‘frontier AI lab’ have the meanings given the terms in section 3 of the Responsible Advanced Artificial Intelligence Act of 2024; and

“(B) the term ‘covered acquisition’ means an acquisition—

“(i) subject to this section; and

“(ii) in which the acquiring person or the person whose voting securities or assets are being acquired is a frontier AI lab.

“(2) Any notification required under subsection (a) for a covered acquisition shall be submitted to the Administration.

“(3) The Administration is authorized to require the submission of additional information or documentary material relevant to a covered acquisition.

“(4) The Administration may submit a recommendation to the Federal Trade Commission and the Assistant Attorney General on whether the covered acquisition violates any of the purposes of the Administration under section 4 of the Responsible Advanced Artificial Intelligence Act of 2024.

“(5) The Federal Trade Commission and the Assistant Attorney General—

“(A) shall cooperate with the Administration in determining whether a covered acquisition, if consummated, would violate the antitrust laws or the purposes of the Administration under section 4 of the Responsible Advanced Artificial Intelligence Act of 2024;

“(B) may use the recommendation of the Administration as a basis for rejecting the covered acquisition or for imposing additional requirements to consummate the acquisition, even if the covered acquisition does not violate the antitrust laws but violates other purposes of the administration under section 4 of the Responsible Advanced Artificial Intelligence Act of 2024; and

“(C) in making a determination described in subparagraph (A), shall give substantial weight to the recommendation of the Administration.”.

SEC. 19. PREEMPTION.

This Act is not intended to preempt any State law, except that any State law or regulation shall be void to the extent that it purports to allow any activity related to frontier AI systems on terms that are less safe or less strict than the terms of this Act. This Act is not intended to preempt any State causes of action, except to the extent that such causes of action directly and substantially interfere with the Administration.

SEC. 20. AUTHORIZATION OF FUNDING.

(a) FROM APPROPRIATIONS.—There are authorized to be appropriated for each fiscal year such sums as are necessary to carry out the purposes of this Act.

(b) FROM LICENSING FEES.—The Administrator may spend fees collected for frontier AI permits as may be necessary to carry out the purposes of this Act.

(c) FROM FINES AND PENALTIES.—The Administrator may spend civil and criminal penalties collected pursuant to this Act as necessary to carry out the purposes of this Act.

(d) FROM DONATIONS.—The Administrator may spend donations received under section 4(d) of this Act.

SEC. 21. SEVERABILITY.

The primary purpose of this Act is to reduce major security risks from frontier AI systems. Moreover, even a short interruption in the enforcement of this Act could allow for catastrophic harm. Therefore, if any portion or application of this Act is found to be unconstitutional, the remainder of the Act shall continue in effect except in so far as this would be counterproductive for the goal of reducing major security risks. Rather than strike a portion of the Act in such a way as to leave the Act ineffective, the Courts should amend that portion of the Act so as to reduce major security risks to the maximum extent permitted by the Constitution.