Responsible 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 Artificial Intelligence Act of 2025".

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

SEC. 2. SENSE OF CONGRESS.

It is the sense of Congress that—

(a) in recent years, artificial intelligence (AI) has rapidly grown more powerful.

(b) computer scientists do not fully understand how advanced AI systems work, nor do they know how to reliably control advanced AI systems;

(c) without additional transparency and oversight measures, we cannot be confident that advanced AI systems will not be used by bad actors to develop bioweapons, launch cyberattacks on critical infrastructure, or otherwise harm national security or public safety

(d) leading AI developers have acknowledged that private AI companies lack the right incentives to fully address this risk; and

(e) it is in the best interests of the United States to track large concentrations of the specialized semiconductors used for advanced AI, require the largest AI developers to adopt commonsense safeguards, and 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 Administration established under section 4 of this Act.

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

(c) ALIGNMENT.—The term "Alignment" means ensuring that AI systems pursue goals that match human values or interests rather than unintended or undesirable goals.

(d) 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).

(e) AI SYSTEM—The term "AI system" means a particular model, program, or tool within the field of AI, or specialized hardware intended for use in developing or operating AI. A collection of AI models, wrappers, plug-ins, and other tools may qualify 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. A collection of AI hardware stored at different locations or owned by different persons may qualify as a single AI system if the hardware is being used in a coordinated fashion to achieve a common purpose.

(f) BOARD.—The term "Board" means the Artificial Intelligence Appeals Board established under section 11 of this Act.

(g) BOARD MEMBER.—The term "Board Member" means a Board Member of the Artificial Intelligence Appeals Board established under section 14 of this Act.

(h) 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.

(i) DEPLOYMENT.—The term "deployment" means to take an AI system that has not yet been made widely available and either—

(1) provide access to that AI system to the general public or to customers, or

(2) make use of that AI system for research, development, internal operations, or any business or commercial purpose, except that using an AI system is not considered deployment if the use is solely for the purpose of (i) evaluating that AI system, or (ii) improving that AI system's safety, alignment, robustness, cybersecurity, or interpretability.

(j) FLOP.—The term "FLOP" means a half-precision (16-bit) floating point operation, which is a measure of compute. Whenever possible, developers must use 16-bit operations for purposes of calculating their compliance with this law. When developers cannot do so, they must instead multiply the number of operations by the bitlength and then divide by 16. The term "FLOPs" is the plural of "FLOP." The term "FLOP/s" means floating point operations per second.

(k) FRONTIER AI.—The terms "Frontier AI" and "Frontier AI system" mean any AI system that:

(1) is a major AI hardware cluster as defined by section 3(s)(9);

(2) is a general-purpose large AI model developed by a medium-compute AI developer that has exhibited dangerous capabilities on automated benchmarks and has been so notified by the Administrator pursuant to section 9(e);

(3) is a general-purpose large AI model developed by a high-compute AI developer;

(4) is ranked as one of the 10 most powerful AI systems in the world based on the benchmarks and weighting published by the Deputy Administrator for Standards; or

(5) has been designated as Frontier AI by the Administrator based on a finding published in the Federal Register and supported by substantial evidence that the AI system poses major security risks that are not otherwise widespread in the AI industry.

(l) GENERAL-PURPOSE AI.—The term "general-purpose AI" means an AI system that is capable of generating outputs with an unrestricted or open-ended format or structure, including but not limited to freeform natural language text, images, audio, video, or code that allows the AI to communicate information, ideas, or instructions in an unconstrained or conversational form or that allows the AI to act as an autonomous or semi-autonomous agent. This includes AI systems designed to engage in general conversation, provide wide-ranging assistance across multiple domains, or answer questions using natural language, even if the system was initially trained for or marketed as supporting a particular subject matter, and even if the system has been fine-tuned or equipped with guardrails that limit the content of its outputs.

(m) 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 close companion living in the household of the individual.

(n) MAJOR SECURITY RISK.—The term "major security risk" includes—

(1) risks that credibly threaten to substantially damage America's public safety, critical 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 operate in a manner that evades or defeats human control.

(o) 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.

(p) NARROW-PURPOSE AI.—The term "narrow-purpose AI" means an AI system that is only capable of producing a particular kind of output with well-understood properties that do not support sustained conversation or interaction with the AI's environment, such as a price, a location, a rating, or a forecast.

(q) QUARTER.—The term "quarter" means January through March, April through June, July through September, or October through December.

(r) 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.

(s) TECHNICAL THRESHOLDS.—The "technical thresholds" in this subsection are initially defined as set forth below, and may be modified and added to by the Administrator and by the Deputy Administrator for Standards pursuant to Section 5(d) and Section 6:

(1) LOW-COMPUTE AI DEVELOPER.—The term "low-compute AI developer" means an AI developer who controlled no more than 10^{18} FLOP/s at any time during the previous quarter.

(2) MEDIUM-COMPUTE AI DEVELOPER.—The term "medium-compute AI developer" means an AI developer who, at any time during the previous quarter, controlled at least 10^{18} FLOP/s, but who at no time during the previous quarter controlled 10^{20} FLOP/s.

(3) HIGH-COMPUTE AI DEVELOPER.—The term "high-compute AI developer" means an AI developer who, at any time during the previous quarter, controlled at least 10^{20} FLOP/s.

(4) LARGE AI MODEL.—The term "large AI model" means an AI system that was trained using at least 10^{26} FLOP. This includes pre-training and fine-tuning. This does not include exploratory runs or general research that did not directly contribute to the AI system's model weights.

(5) LARGE TRAINING RUN.—The term "large training run" means any training run that is intended or expected to result in the creation of a large AI model.

(6) SIGNIFICANT DEPLOYMENT.—The term "significant deployment" means deploying an AI system so as to (i) perform inferences that generate more than 1 gigabyte of output or more than 100 million tokens of output, (ii) make the AI system or its model weights generally available to the public, or (iii) support or conduct recursive self-improvement.

(7) HIGH-PERFORMANCE AI CHIP—The term "high-performance AI chip" means any integrated circuit covered by ECCN 3A090.a 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.

(8) SIGNIFICANT AI HARDWARE TRADERS.—The term "significant AI hardware trader" means a person who, during the immediate previous quarter, bought, sold, gifted, received, rented, traded, destroyed, or transported at least 100 high-performance AI chips. (For this definition, a person "transports" a chip if they cause its location to change by at least 10 miles.)

(9) MAJOR AI HARDWARE CLUSTER—The term "major AI hardware center" means at least 3,000 high-performance AI chips that (A) are being used or made available for a common or coordinated purpose, and (B) are either owned by the same person, or are substantially controlled by a joint venture, consortium, partnership, or other organized project that is able to direct how the chips shall be used. Chips stored at different locations may nevertheless be part of the same major AI hardware cluster.

(t) TRAINING.—The term "Training" means the process of fitting model weights 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 ADMINISTRATION.

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

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

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

(b) MISSION.—It shall be the mission of the Frontier Artificial Intelligence Administration to mitigate major security risks by monitoring and analyzing the most important threats posed by the largest and most advanced AI systems, by developing common-sense safeguards and standards for AI systems to follow in order to avoid these threats, and by enforcing these standards and requiring these safeguards.

(c) 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.

(d) 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 one year 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 one year 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.

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

(A) LIMIT ON QUANTITY.—The Administrator may not employ more than 100 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.

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

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

(1) Part III of title 5, United States Code or section 1342 of title 31, United States Code shall not bar such volunteers from service.

(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.

(g) 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.

(h) INFORMATION SECURITY.

(1) The Administrator shall, not later than 180 days after the enactment of this Act, establish and maintain information security protocols for the Administration to ensure the secure handling of sensitive information acquired under this Act, including information related to permit applications, evaluations, and frontier AI systems. These protocols shall—

(A) meet or exceed Federal Information Security Management Act requirements;

(B) provide appropriate protection for classified information and controlled unclassified information;

(C) ensure secure handling of sensitive information about frontier AI systems and capabilities;

(D) include an incident response and mitigation plan for breaches or unauthorized disclosures of sensitive information.

(2) The Administrator shall review and update these protocols annually to address evolving security threats.

(3) The Administrator may designate certain positions within the Administration as requiring security clearances based on access to sensitive information.

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) HOW REMOVED.—The Administrator may remove any Deputy Administrator for any or no reason, except that each such removal must be published in the Federal Register as described in Section 15(f)(2), and the Deputy Administrator for Public Interest cannot be removed in such a way as to interfere with the protections in Section 5(h)(2) without a direct written order from the President.

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

(d) 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 15 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 chips, and shall organize efforts to detect and identify any high-performance AI chips that have 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.

(e) DEPUTY ADMINISTRATOR FOR STANDARDS.—The Deputy Administrator for Standards shall be responsible for supervising and directing the progress of rulemaking under Sections 6 and 11 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 technical thresholds in Section 3(s), accounting for gains in algorithmic efficiency;

(4) within 60 days after being appointed, select and publish in the Federal Register a set of one or more benchmarks to be used to identify frontier AI systems as described in Section 3(k)(4), these benchmarks shall be selected from among 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;

(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, and

(6) within 90 days after being appointed, either select and publish in the Federal Register a set of one or more automated benchmarks to automatically detect dangerous capabilities in large AI models created by medium-compute AI developers as described in Section 9, or, if no such benchmarks are available, describe in the Federal Register the reasons why existing benchmarks are inadequate and issue a call for proposals for improved benchmarks for this purpose. The call for proposals would be supported by grants as described in Section 5(f)(iv) to then facilitate the development of such benchmarks.

(f) 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.

(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:

(1) 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.

(2) 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 chips or that render such hardware easier to remotely monitor or remotely disable.

(3) SOFTWARE SAFETY RESEARCH.—Funding research, development, and prototypes of software safety features for AI systems, especially mechanistic interpretability, alternative AI architectures that are fundamentally secure by design, and verifiable corrigibility.

(4) 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. If the Administrator issues a call for proposals for improved benchmarks pursuant to Section 5(d)(6), then such proposals shall receive the highest priority under this paragraph.

(5) RESEARCH INTO POST-DEPLOYMENT RESILIENCE.—Funding research, development, and prototypes of failsafes, backups, incident trackers, kill switches, antidotes, firewalls, analog tools for critical infrastructure, and other resources that increase social resilience to the major security risks that may arise from frontier AI systems.

(6) 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.

(B) 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.

(C) 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.

(3) DISMISSAL BY THE PRESIDENT.—Notwithstanding the additional protections in the paragraph above, the President may dismiss the Deputy Administrator for Public Interest at any time for any reason. Any such dismissal shall be in writing and shall be promptly published in the Federal Register. After such a dismissal, the Administrator shall promptly appoint a replacement as described in section 5(a).

(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 18 of this Act.

SEC. 6. RULEMAKING AUTHORITY.

(a) IN GENERAL.—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 thresholds in Section 3(s) of this Act (including but not limited to the definitions of "high-compute AI developer," "high-performance AI chip," and "major AI hardware cluster") 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.

(b) STANDARD FOR ALTERING DEFINITIONS.—Before modifying any technical threshold in this Act, the Administrator must first publish findings in the Federal Register supported by

clear and convincing evidence that the proposed modifications will not significantly increase major security risks.

(c) TECHNICAL THRESHOLDS FOR BIOLOGICAL, CHEMICAL, AND OTHER SPECIALIZED DATA.—The Administrator is encouraged to develop and promulgate, at the earliest feasible time, technical thresholds to identify AI systems that qualify as frontier AI based on having trained on specialized data related to biology, chemistry, weapons, or other technical subjects that may yield capabilities that are significantly more powerful compared to the capabilities that typically arise from training on a comparable amount of ordinary data.

(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) 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.

SEC. 7. REPORTING FOR SIGNIFICANT AI TRANSACTIONS.

(a) REPORTING REQUIREMENT.—Within 30 days of the start of each quarter, each significant AI hardware trader must submit a report to the Administrator indicating all of the following information:

(1) the total number of high-performance AI chips they acquired;

(2) the total number of high-performance AI chips that they sold or otherwise disposed of; and

(3) all of the addresses at which at least 10 of the high-performance AI chips they own are currently located.

(b) WEBSITE FOR COLLECTING REPORTS.—No later than 90 days after this law is enacted, the Administrator shall create a website containing a form to be filled out by significant AI hardware traders to submit the information required by this section. The Administrator shall advertise the availability of the website 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.

(c) WHEN OBLIGATION BEGINS.—Significant AI hardware traders must begin making the reports required by this section during the first quarter that begins at least 60 days after the notice of the reporting website's availability is published in the Federal Register.

SEC. 8. HARDWARE SECURITY AUDITS.

(a) WHEN HARDWARE LICENSE NEEDED.—A person must seek and receive a hardware license before owning, importing, leasing, renting, controlling, or knowingly possessing a major AI hardware cluster.

(b) QUALIFIED INDEPENDENT AUDITOR.—In order to apply for a hardware license, an applicant must first receive a cybersecurity audit from a qualified independent contractor. The audit must meet all of the following criteria:

(1) The auditor's compensation is not connected to the results of the auditor's findings.

(2) No more than one-third of the auditor's annual revenue is derived from or subject to the approval of the same audit recipient.

(3) The auditor has the technical expertise necessary to competently evaluate the applicant's cybersecurity, physical security, and customer verification plan.

(4) The auditor receives adequate access to the applicant's data, plans, facilities, and personnel so as to be able to conduct a thorough and open-ended investigation.

(5) The auditor receives adequate time to conduct its investigation.

(6) The auditor receives adequate time to write up its report after the investigation.

(c) CONTENTS OF AUDIT REPORT.—After completing its investigation, the auditor shall express an explicit opinion as to each of the following:

(1) The auditor's technical qualifications and financial independence.

(2) The adequacy of the auditor's access to the applicant's data, plans, facilities, and personnel.

(3) The adequacy of the time provided to the auditor for the investigation and report.

(4) Whether the applicant's major AI hardware cluster will be reasonably secure against cyberattacks, especially cyberattacks that could allow unauthorized third parties to access a significant amount of the applicant's compute or to exfiltrate model weights from the cluster.

(5) Whether the applicant's major AI hardware cluster will be reasonably secure against physical theft, especially theft that could result in an inability to account for or retrieve high-performance AI chips.

(6) Whether the applicant will be able to reliably determine the real identities of its major AI hardware cluster's major customers, if any.

(d) SUBMISSION OF AUDIT REPORT.—The auditor shall submit their audit report directly and privately to the Administrator, at least 30 days before the applicant receives access to or copies of the audit report. The Administrator will then evaluate the audit report. The Administrator may reject the hardware license application if and only if at least one of the following conditions is met:

(1) the audit failed to meet one of the requirements of Section 8(b);

(2) the audit report failed to meet one of the requirements of Section 8(c);

(3) the auditor did not offer an opinion based on reasonable assurance that the major AI hardware cluster will not significantly contribute to major security risks;

(4) the audit report placed significant reliance on patently unreasonable assumptions;

(5) the audit report placed significant reliance on false or fraudulent data;

(6) the reasoning or conclusions of the audit report were arbitrary or capricious; or

(7) the Administrator can demonstrate that approving the application would severely exacerbate major security risks.

(e) SCOPE OF HARDWARE LICENSE.—After receiving a hardware license, a person may improve, expand, or reconfigure their major AI hardware clusters. However, such a person must apply to renew their license:

(1) at least annually, i.e., no more than one year after the person's most recent license was granted;

(2) each time the number of high-performance AI chips under the person's control increases by a factor of 10 or more;

(3) if the person is a privately held corporation, each time there are substantial changes to the person's ownership; and

(4) before weakening or cancelling any policies or procedures used to maintain cybersecurity, maintain physical security, or maintain awareness of the real identities of the person's customers.

SEC. 9. AUTOMATED BENCHMARKING FOR MEDIUM-COMPUTE AI DEVELOPERS.

(a) SCOPE OF REQUIREMENT.—Before a medium-compute AI developer makes a significant deployment of a large AI model, that developer shall conduct automated benchmark testing on the large AI model.

(b) WHICH BENCHMARKS TO USE.—The benchmarks to be used will be selected, maintained, and made available by the Administrator based on their ability to detect

dangerous capabilities and on their ease of use. To be included as a required benchmark under the section, the benchmark must meet all of the following criteria:

(1) it is privacy-preserving, i.e., completion of the benchmark does not require sharing the AI developer's model weights or algorithms;

(2) it is substantially automatic, i.e., a typical medium-compute AI developer must be able to submit a data file to a website or other portal and receive a result within 24 hours;

(3) it is reasonably objective, i.e., the benchmark yields scores that can be meaningfully compared across AI systems and that do not depend on the personal preferences of a human grader.

(c) REPORT ON BENCHMARK RESULTS.—After completing the automated benchmarks, a medium-compute AI developer shall submit a report to the Administrator containing all of the following information:

(1) The name and contact information of the person responsible for the large AI model's training.

(2) The amount of compute used during the large AI model's training run.

(3) The date on which the large AI model's training run was completed.

(4) The general purpose of the large AI model.

(5) The principal addresses at which the large AI model was trained, or, if the large AI model was trained using cloud computing resources, the name of the principal cloud computing providers used for the training.

(6) The scores received by the large AI model on the automated benchmarks.

(d) WHEN TO BE COMPLETED.—The report must be completed and submitted to the Administrator within 45 days after the significant deployment of the large AI model.

(e) WHEN OBLIGATION BEGINS.—A medium-compute developer is not required to submit any reports based on large AI models that were significantly deployed no later than 30 days after the Administrator first publishes the list of benchmarks to be used.

(e) DISCOVERY OF DANGEROUS CAPABILITIES.—If, based on the report received pursuant to this section, the Administrator determines that a large AI model poses major security risks, then the Administrator shall so notify the medium-compute AI developer in writing. A medium-compute AI developer that has received such a notice must promptly take all of the following actions:

(1) to the extent technically feasible, pause or cancel all commercial access to the large AI model;

(2) refrain from significantly deploying any other large AI models;

(3) refrain from selling or transferring the model weights or algorithms associated with the large AI model; and

(4) take all appropriate precautions to prevent the leak, unauthorized publication, or exfiltration of model weights or algorithms associated with the large AI model.

(f) APPLICATION FOR HIGH-COMPUTE DEVELOPER SOFTWARE PERMIT.—After receiving a notice that its AI system poses major security risks as described in Section 9(e), a medium-compute developer may apply for one or more software permits as set forth in Section 10. The medium-compute developer may resume deployment of AI systems based on large training runs only after receiving a valid software permit for each such AI system.

SEC. 10. SOFTWARE PERMITS FOR HIGH-COMPUTE AI DEVELOPERS.

(a) SCOPE OF REQUIREMENT.—Each developer who makes a significant deployment of frontier AI software must first apply for and receive a software permit.

(b) EXEMPTIONS.—A developer does not need a software permit if they meet any of the following criteria for an exemption:

(1) the AI system is narrow-purpose, well-understood, or otherwise extremely unlikely to pose major security risks;

(2) the developer controlled less than 10^{20} FLOP/s at all times during the previous quarter; or

(3) the developer used less than 10^{26} FLOP to train the AI system.

(c) TIMING OF REQUIREMENT.

(1) PAST ACTIVITIES.—A person who has already deployed an AI system as of the date this law is enacted may indefinitely continue owning or operating that AI system 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 high-compute developer may not make a significant deployment of a new large AI model without a permit simply because they deployed a similar large AI model before this law was passed.

(2) ONGOING ACTIVITIES.—A person who needs a software permit based on ongoing activities that began prior to the date when the Administrator first publishes the relevant application form 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 software permit based on activities they are conducting after the Administrator first publishes the relevant application form must apply for and receive that permit before beginning those activities.

(d) SOFTWARE AUDIT. In order to apply for a software permit, a developer must first receive a software audit from a qualified independent contractor. The audit must meet all of the following criteria:

(1) The auditor's compensation is not connected to the results of the auditor's findings.

(2) No more than one-third of the auditor's annual revenue is derived from or subject to the approval of the same audit recipient.

(3) The auditor has the technical expertise necessary to competently identify and evaluate any major security risks posed by the applicant's proposed AI systems.

(4) The auditor has the technical expertise necessary to competently evaluate the extent to which the applicant's proposed guardrails would adequately mitigate any major security risks posed by the applicant's proposed AI systems.

(5) The auditor receives adequate access to the applicant's data, plans, facilities, and personnel so as to be able to conduct a thorough and open-ended investigation.

(6) The auditor receives adequate time to conduct its investigation.

(7) The auditor receives adequate time to write up its report after the investigation.

(e) CONTENTS OF AUDIT REPORT.—After completing its investigation, the auditor shall express an explicit opinion as to each of the following:

(1) The auditor's technical qualifications and financial independence.

(2) The adequacy of the auditor's access to the applicant's data, plans, facilities, and personnel.

(3) The adequacy of the time provided to the auditor for the investigation and report.

(4) What major security risks, if any, are posed by the applicant's proposed AI systems, how severe these risks would be if they occurred, and how likely they are to occur.

(5) What guardrails, if any, are proposed by the applicant to mitigate each major security risk posed by the applicant's proposed AI systems.

(6) Whether the proposed guardrails adequately mitigate the major security risks, if any, posed by the applicant's proposed AI systems, and, if so, why.

(f) SUBMISSION OF AUDIT REPORT.—The auditor shall submit their audit report directly and privately to the Administrator, at least 30 days before the applicant receives access to or copies of the audit report. The Administrator will then evaluate the audit report. The Administrator may reject the software permit application if and only if at least one of the following conditions is met:

(1) the audit failed to meet one of the requirements of Section 10(d);

(2) the audit report failed to meet one of the requirements of Section 10(e);

(3) the auditor did not offer an opinion based on reasonable assurance that deploying the AI system will not significantly contribute to major security risks;

(4) the audit report placed significant reliance on patently unreasonable assumptions;

(5) the audit report placed significant reliance on false or fraudulent data;

(6) the reasoning or conclusions of the audit report were arbitrary or capricious; or

(7) the Administrator can demonstrate that approving the application would severely exacerbate major security risks.

SEC. 11. DEVELOPMENT OF APPLICATION FORMS

(a) "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.

(1) 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.

(2) 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.

(3) 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.

(4) 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—

(A) self-driving cars;

(B) navigational systems;

(C) recommendation engines;

(D) fraud detection systems;

(E) weather forecasting tools;

(F) tools for locating deposits of oil, gas, or minerals;

(G) AI systems designed to predict the demand, supply, price, cost, or transportation needs of products or services;

(H) search engines whose primary function is to suggest webpages; and

(I) AI systems whose function is substantially limited to generating still images, each of which typically contains no more than thirty words of text.

(b) INITIAL APPLICATION FORMS.—No later than six months following the enactment of this Act, the Administrator shall promulgate an initial application form for hardware permits and an initial application form for software permits.

(c) RENEWAL APPLICATION FORMS.—No later than twelve months following the enactment of this Act, the Administrator shall promulgate a renewal application form for hardware permits and a renewal application form for software permits. Each permit shall expire and require renewal 1 year after it is issued, unless the Administrator promulgates a rule varying this time period. A person who makes substantial changes to the design, use cases, or safeguards related to their frontier AI system shall submit a renewal application for the changed system within 30 days of making the relevant changes. There is a rebuttable presumption that an AI system that appears in any press release or advertisement with a new name or a new version number has undergone substantial changes.

(d) 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.

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

(f) 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.

(1) 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.

(2) 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.

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

(4) 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.

(g) RUBRICS FOR APPLICATIONS.—No later than eight months following the enactment of this Act, the Administrator shall promulgate advisory scoring rubrics that auditors are encouraged to use in order to evaluate applications for AI hardware and software permits.

(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 to correspond to a recommendation that an application be approved.

(2) MISSING CRITERIA.—In developing each rubric, the Administrator shall consider assigning scoring factors based on each of the criteria listed in Section 12. For each such criterion that is not incorporated into a scoring factor on the relevant rubric, the Administrator shall issue a written statement in the Administration's annual bulletin in the Federal Register explaining why the criterion was not included.

(3) 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 or remove one of the scoring factors recommended by Section 12 unless the Administrator first publishes an explanation in the Federal Register at least 45 days before that update takes effect.

(4) USE OF RUBRICS.—An independent auditor shall not be required to use the rubrics recommended by the Administrator, but the Administrator may consider the extent and justifiability of an auditor's deviation from the rubrics when evaluating whether an audit report was arbitrary or incomplete.

SEC. 12. RECOMMENDED SCORING FACTORS

(a) RUBRICS FOR SOFTWARE PERMITS.—The Administrator should recommend scoring factors for software audits based on each of the following criteria:

(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 applicant's previous track record of accurately forecasting the capabilities and risks of their advanced AI systems;

(8) the applicant's plan for reserving a substantial and specific fraction of the total compute budget for the purpose of safety testing and safety research;

(9) the applicant's plan for ensuring that the AI system it is training will not escape during training, e.g., by being copied to systems outside of the applicant's control, or by

operating on systems other than those intended by the applicant, or otherwise significantly influencing the world outside of the laboratory in which it is being trained;

(10) the applicant's plan for which specific persons or job roles will receive access to each type of data, each type of algorithm, and each set of model weights during training, and for how unauthorized persons will be denied access to such information;

(11) 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.

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

(13) the extent to which the applicant's AI systems could autonomously survive, replicate, or spread;

(14) the extent to which the applicant's AI systems directly contributes to activities such as bioweapons development, nuclear weapons development, or automated hacking;

(15) the applicant's resources, abilities, reputation, and willingness to successfully execute the plans described in the other paragraphs in this subsection; and

(16) the extent to which the applicant has provided complete, accurate, and timely information about its AI systems, including proactive disclosures of potential sources of major security risks, to its auditors and to other appropriate parties.

(b) DEGREE OF OPENNESS.—An applicant for a software permit shall clearly indicate in their application the extent to which the applicant plans to share access to the resulting algorithms and/or model weights. The applicant shall indicate who will be allowed to access this information, what actions they must perform in order to gain access, and whether the access will be gated or conditioned in any way, such as via an API.

(c) ADDITIONAL SCORING FACTORS FOR CLOSED-SOURCE AI SYSTEMS.—For applicants who indicate that they intend to restrict access to their AI system to any significant extent, the Administrator should also include scoring factors for each of the following criteria:

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

(2) the applicant's plan for ensuring that its advanced 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;

(3) the applicant's plan for ensuring that its advanced AI system will not be shared with unauthorized users;

(4) 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;

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

(6) the applicant's resources, abilities, reputation, and willingness to successfully execute the plans described in the other paragraphs in this subsection.

(d) ADDITIONAL SCORING FACTORS FOR OPEN-SOURCE AI SYSTEMS.—For applicants who indicate that they intend to provide public access to their AI system to any significant extent, the Administrator should also include scoring factors for each of the following criteria:

(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 that its AI system will remain safe even after being fine-tuned, connected with plug-ins or utilities, or otherwise modified by users who have received access to the AI system's source code or model weights;

(3) the applicant's analysis of how the AI system might acquire new capabilities in the years after it is released, and their analysis of why these new capabilities will not pose major security risks;

(4) the applicant's plan for ensuring the traceability of products and services enabled by the applicant's AI system;

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

(6) the applicant's resources, abilities, reputation, and willingness to successfully execute the plans described in the other paragraphs in this subsection.

(e) RUBRICS FOR HARDWARE PERMITS.—The Administrator should recommend scoring factors for hardware security audits based on each of the following criteria:

(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;

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

(4) the applicant's plan for ensuring that small businesses and entrepreneurs have the first option to purchase at least 10% of the hardware resources, if any, that the applicant rents or otherwise makes available to customers or clients.

SEC. 13. ADJUDICATION OF PERMIT APPLICATIONS.

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

(b) INITIAL REFERRAL.—Upon receipt of an application for a hardware permit or software permit under this Act, the Administrator shall refer the application to two randomly selected

AIPJs and shall also forward a copy of the application to the Deputy Administrator for Public Interest.

(c) OBJECTION BY AIPJS.—Within 60 days after receipt of an application, either of the AIPJs assigned to review that application may object to the application on one or more of the following grounds:

(1) the auditor was unqualified;

(2) the audit report was significantly incomplete;

(3) the audit report placed significant reliance on patently unreasonable assumptions;

(4) the audit report placed significant reliance on false or fraudulent data;

(5) the reasoning or conclusions of the audit report were arbitrary or capricious; or

(6) approving the application would severely exacerbate major security risks.

(d) OBJECTION BY DEPUTY ADMINISTRATOR.—Within 60 days after receipt of an application, the Deputy Administrator for Public Interest may object to the application on the grounds that it would severely exacerbate major security risks.

(e) FORMAT OF OBJECTIONS.—An objection registered under this section shall be in writing, and shall be signed and dated by the person making the objection, and shall clearly explain why the application is objectionable.

(f) RECOMMENDATION OF APPROVAL BY DEFAULT.—If after 60 days, neither the AIPJs nor the Deputy Administrator have objected to the application, then the application will be considered to be recommended for approval by default. In that case, the Administrator shall promptly either accept the recommendation and notify the applicant that their application was approved, or initiate an appeal pursuant to Section 14.

(h) CONFERENCE FOLLOWING OBJECTION.—If either of the AIPJs or the Deputy Administrator objects to the application, then the Deputy Administrator for Public Interest or their designee shall meet and confer with the two AIPJs and attempt to reach a consensus on how to resolve the objection. If they cannot reach agreement within 30 days after the objection is registered, then the application is considered to have been recommended for rejection. Alternatively, by agreement of at least two of those three persons, the reviewers may take any one of the following actions with respect to the 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; or

(4) recommend rejection of the application.

(h) OPINION BY AIPJ.—An AIPJ who formed part of the majority for each recommendation shall provide a short written opinion explaining the basis for the AIPJ'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.

(i) FINALITY OF RECOMMENDATIONS.—A recommendation becomes final immediately if it is approved by the Administrator, the Deputy Administrator for Public Interest, and the apploicant. A recommendation also 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.

(j) POWER TO REQUIRE PRECAUTIONS.—An application that is approved with conditions may include one or more of the following conditions:

(1) The applicant must 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 cyberattacks.

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

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

(4) The applicant must apply for and receive a certain type and amount of insurance coverage before conducting activities under the permit.

(5) Other precautions that the Administrator finds to be useful or appropriate.

(k) CONSIDERATIONS FOR OPEN SOURCE FRONTIER AI SYSTEMS.

(1) COSTS AND BENEFITS.——When evaluating an application to deploy an open-source AI system, all evaluators shall 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.

(2) OPTIONS.—The Administrator may reject an application for an open-source frontier AI system or may approve an application for an open-source frontier AI system with or without conditions.

(3) TYPES OF LIMITATIONS.— 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.

(4) TARGETS OF LIMITATIONS.—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.

(5) NO AUTOMATIC DETERMINATIONS.—An 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 10 through 14 and shall determine whether the accompanying software audit satisfactorily demonstrates that training the AI system and deploying it on an open-source basis will not severely exacerbate major security risks.

(6) NO APPLICATION TO NON-FRONTIER OPEN SOURCE.—The guidance in paragraphs (1) through (3) above applies only to frontier AI models. AI models that are narrow-purpose, based on a small training run, or conducted by a low-compute AI developer are exempt from the auditing and licensing process.

SEC. 14. APPEALS OF PERMIT APPLICATIONS.

(a) TIMING AND PARTIES FOR APPEAL.—The applicant, the Deputy Administrator for Public Interest, or the Administrator may appeal an AIPJ's recommendation to the AI Appeals Board by filing a notice of appeal within 20 days of receiving the AIPJ's written opinion. The notice of appeal shall consist of a short, plain statement of the facts and reasoning supporting the appeal.

(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, without being bound by the legal reasoning of the AI Permit Judges. 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 AIPJs, at the discretion of the Appeals Board. The Appeals Board shall resolve each appeal before it and any procedural questions associated with that appeal 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 section 14(d), 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 14(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 permit 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. 15. 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 8 and attempt to identify and describe all of the following:

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

(B) The most notable concentrations of high-performance AI chips, especially collections of large numbers of high-performance AI chips 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 chips.

(D) Notable changes in the flow of high-performance AI chips 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 chips.

(C) The rate at which high-performance AI chips are 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 chips are being used.

(3) No later than the 20th day of each month, the Administrator shall collate and analyze the information collected via paragraphs (1) and (2) with the goal of determining which high-performance AI chips (if any) are 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, obtain judicial warrants permitting entry onto premises where frontier AI systems are reasonably believed to exist, and require production 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) REVIEW OF THRESHOLDS.—No later than September 1st of each year, the Administrator shall review each relevant threshold and technical definition in this Act, and determine whether each threshold and each technical definition remains adequate to defend against major security risks. If any threshold or technical definition has become inadequate, then the Administrator shall promptly promulgate rules to appropriately strengthen or tighten the threshold or technical definition. The Administrator is not required to review these thresholds or technical definitions during the same calendar year that the Act is enacted.

(f) 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(e)(2) of this Act.

(2) The names and positions of Deputy Administrators who have been removed, the dates on which they were removed, and, for each such removal, a description of the cause for which the Deputy Administrator was removed from their position, or a statement that the Deputy Administrator was removed without 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 permits 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 18, 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. 16. CIVIL LIABILITY.

(a) WHO OWES DUTY OF CARE.—All persons engaged in the development or deployment of frontier AI systems 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–

(1) knowingly owns, controls, creates, or makes a significant deployment of a frontier AI system,

(2) publishes, sells, leaks, or transmits the model weights of a frontier AI system to any third party, or

(3) owns, controls, possesses, or assembles a major AI hardware cluster.

(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 frontier 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 unintentionally available to the public.

(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 of their frontier AI or credible threat to misuse their frontier AI, immediately take all practical steps to deny that third party access to their frontier AI.

(5) Their major AI hardware clusters 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 major AI hardware clusters,

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

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

(d) JOINT AND SEVERAL LIABILITY.—All persons who have violated the duty of care imposed by this Act with respect to the same frontier AI system 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 $1 billion in tangible damages based on a violation of the duty of care established by 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 $1 billion 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.

(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 $250,000 and no more than $5,000,000 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 after receiving a judgment or preliminary injunction based on that violation may be assessed an additional civil penalty of up to $1,000,000 for each day that the violation continues after receipt of the judgment or preliminary injunction.

(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) *PER SE* VIOLATIONS.—In any lawsuit alleging a violation of the duty of care created by this subsection, a defendant shall be deemed to have committed a *per se* violation of that duty of care 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 misrepresentation (including a misrepresentation by omission) to their auditors or to the government as part of their application for a frontier AI permit, and this misrepresentation was material in the sense that it (i) could plausibly have affected the auditor's recommendation or the government's decision, and (ii) could plausibly have caused or exacerbated a harm that is a subject of the lawsuit.

(h) STRICT LIABILITY.—In any civil lawsuit where the plaintiff or plaintiffs demonstrate by a preponderance of the evidence that a defendant committed a *per se* violation of their duty of care, that 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) 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.

(k) 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.

(l) SAFE HARBOR FOR SMALL BUSINESSES AND ORGANIZATIONS.—It shall be a complete affirmative defense to any civil lawsuit brought under this section that, at all times covered by the allegations of the complaint, a defendant met all of the following criteria:

(1) the defendant employed fewer than 500 full-time equivalent persons;

(2) the defendant controlled less than $500 million in assets;

(3) the defendant had a valid frontier AI hardware permit covering all major AI hardware clusters under its control, if any;

(4) the defendant had a valid frontier AI software permit covering all frontier AI models under its control, if any;

(5) the defendant did not commit fraud in obtaining any frontier AI permit; and

(6) the defendant substantially complied with the conditions, if any, imposed by all of its frontier AI permits.

SEC. 17. 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 18 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.

(6) The person meets all of the following criteria:

(A) the person knows that they are required to apply for and receive a frontier AI permit before conducting a particular activity, which may be demonstrated *prima facie* by showing that the person has been sent a letter from the Administration informing them that an AI system related to that activity qualifies as frontier AI;

(B) the person knowingly conducts that particular activity, and

(C) the person knows that they have not yet received a frontier AI permit that would cover that particular activity.

(7) The person improperly uses a shell corporation or other legal fiction with the intent of misleading the government as to the person's total compute assets or total compute budget so as to wrongfully evade one of the requirements of this Act.

(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 (i) knowingly acquires a major AI hardware cluster or knowingly makes a significant deployment of a frontier AI system, and (ii) such activity requires a frontier AI permit, and (iii) the person does not have a currently valid frontier AI permit that covers this activity.

(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.

(5) The person knowingly forms, operates, or controls a corporation that is or will be a high-compute AI developer, and recklessly fails to adequately capitalize that corporation, thereby causing that corporation to be unable to fully discharge liabilities arising out of the corporation's AI-related activities.

(d) SELF-REPORTING INFRACTIONS.—Any significant AI hardware trader who is required to file a quarterly report on transactions involving high-performing AI chips 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 $50,000 or twice the total price of the high-performance AI chips, 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 $1 billion 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, except research that is directly related to and necessary for correcting a problem that contributed to the misdemeanor.

(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 closed-source 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 software permit. The entity shall sell or transfer all of its frontier AI hardware to other person(s) with a valid frontier AI hardware 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. 18. 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. Repeating an executive declaration of emergency powers based on the same or substantially similar situation shall not be effective to renew the term of those powers.

(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) PURPOSE OF EMERGENCY POWERS.—Any emergency powers made available under this section shall be used solely for the purpose of mitigating major security risks from frontier AI.

(f) 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.

(g) 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.

(h) 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.

(i) 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 14 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.

(j) 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,

(2) the order was made without adequate legal authority, or

(3) the order was applied without a reasonable relationship to mitigating major security risks from frontier AI.

(k) 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. 19. 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) testified, assisted, reported, made allegations in, or otherwise participated in any investigation, litigation, hearing, or proceeding directly related to frontier AI;

(2) reported the existence or likelihood of any major security risk from frontier AI to an appropriate superior within the whistleblower's organization;

(3) 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; or

(4) identified a policy, action, or plan as being forbidden by this Act.

(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 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. 20. 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) High-compute AI developers

"(1) In this subsection—

"(A) the terms 'Administration' and high-compute AI developer' have the meanings given the terms in section 3 of the Responsible Artificial Intelligence Act of 2025; 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 high-compute AI developer.

"(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 would be likely to increase major security risks or otherwise conflict with the purposes of the Responsible Artificial Intelligence Act of 2025.

"(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 Responsible Artificial Intelligence Act of 2025;

"(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 Responsible Artificial Intelligence Act of 2025; and

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

SEC. 21. 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. 22. 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. 23. AI SAFETY AND SECURITY FUND.

(a) ESTABLISHMENT.—There is established in the Treasury of the United States a fund, to be known as the "AI Safety and Security Fund" (referred to in this section as the "Fund"), to be administered by the Frontier Artificial Intelligence Administration.

(b) AVAILABILITY.—Amounts deposited in the Fund shall be available to the Administrator, without further appropriation, for any of the following purposes authorized by this Act, including:

(1) Administering and enforcing this Act, including—

(A) hiring personnel,

(B) conducting audits and investigations,

(C) acquiring computing and analytic resources,

(D) maintaining cybersecurity and physical security infrastructure.

(2) Supporting research, grants, and public goods authorized under this Act, including—

(A) development of automated benchmarks and safety evaluations,

(B) interpretability and robustness research,

(C) grants to small businesses and voluntary auditors.

(3) Providing technical assistance and outreach to small developers, minority-serving institutions, and civil society organizations engaged in AI safety, fairness, or public-interest work.

(c) SOURCES OF FUNDS.—The following shall be credited to the Fund:

(1) All civil and criminal fines collected under this Act;

(2) All permit application fees collected under this Act;

(3) All donations and bequests made under Section 4(g) of this Act, unless otherwise restricted by the donor;

(4) Any other funds appropriated by Congress for the purposes described in subsection (b).

(d) FUND MANAGEMENT.—

(1) The Administrator shall ensure that all expenditures from the Fund are publicly reported in the annual bulletin required under Section 15(f).

(2) The Administrator may not obligate more than 10 percent of the Fund's balance in any single fiscal quarter without the written concurrence of at least two Deputy Administrators, one of whom shall be the Deputy Administrator for Public Interest.

(e) LIMITATION.—No funds from the AI Safety and Security Fund may be used to compensate any person in violation of Section 4(d) of this Act, nor to settle any claims for misconduct or negligence by Administration employees.

SEC. 24. 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.