The Responsible AI Act of 2025:
Section-by-Section

Overview

The Center for AI Policy (CAIP) is a nonpartisan research organization that develops policy and conducts advocacy to mitigate catastrophic risks from AI. AI poses a wide variety of important problems, including job loss, invasion of privacy, bias, hallucinations, and threats to national security. Our staff (and this bill) are focused on the most severe and deadly threats that could arise from advanced AI. We hope that Congress will take action to mitigate **all** of the risks from AI, but our particular expertise is in catastrophic risks, so that is what we are writing about.

We are sharing the Responsible Artificial Intelligence Act of 2025 ("RAIA") as model legislation – we hope that Congressional offices will find some or all of it useful as they develop solutions to the problems posed by the rapid advance of AI. Although the various sections of this bill do complement each other, many sections would also function well as standalone bills. Please contact CAIP at info@aipolicy.us for assistance in writing a standalone bill based on one or two sections of RAIA or otherwise adapting this work to meet your office's needs.

Structure of the Bill

Sections 1 through 3 of RAIA define "frontier artificial intelligence systems" and the "major security risks" that frontier AI might soon pose. A frontier AI system is software at the cutting edge of what broadly capable AI can accomplish. Typically, to qualify as frontier AI, an AI system must be highly advanced, general-purpose, and trained by a major developer who is spending hundreds of millions of dollars. Most of the AIs developed by startups, small businesses, nonprofits, and academic researchers will be exempt from this law. Similarly, AIs that can only offer a narrow range of outputs (like a location, a price, or a rating) would also be exempt, because such AIs will not manufacture WMDs or escape from human control.

Section 4 of the RAIA sets up a new federal office called the Frontier AI Administration ("FAIA") that would have a mandate to tackle the unique threats posed by advanced, general-purpose AI. We imagine that the office would have a staff of a few hundred people. As written, the office would be independent (like NASA or the CIA), but it could also be structured to report to the Secretary of Commerce, Energy, or Homeland Security. ***Section 5*** of the bill explains how the office's leadership would be appointed and what the major divisions of the office would do. ***Section 6*** explains how the office would set and update rules to keep pace with changes in technology.

Sections 7 and 8 set minimum standards for advanced AI hardware security. Buyers and sellers who trade more than 100 high-performance AI chips in the same quarter must fill out a short report on their transactions, and owners who run a high-performance computing cluster with at least 3,000 such chips must pass annual audits showing that they maintain adequate cybersecurity protocols and are aware of the real identities of their major customers.

Sections 9 through 14 set minimum standards for advanced AI software security. Major AI developers who control at least 10^{20} FLOP/s (approximately five billion dollars' worth of chips) are required to pass a third-party security audit before deploying large new models. The private auditors will check to see what major security risks might be posed by the model and evaluate whether those risks have been adequately mitigated. The federal AI security office created in Section 4 will then check the audit to make sure it is adequate and complete; if the audit or the audit review reveal evidence that a new AI system would be extremely dangerous, then its developer could be required to add additional safeguards, delay the release of the model, or, if necessary, come up with a new design. Developers who disagree with these requirements will have a right to appeal.

Section 15 tasks the new federal AI office with tracking and monitoring concentrations of advanced AI hardware and software and identifying trends in the movement of these strategic resources. The office will then report on what it learns (as appropriate) to other parts of the government and to the public.

Sections 16 and 17 lay out a civil and criminal liability regime for catastrophic risks from AI. The civil liability would be reserved for suits brought by the federal government based on genuine threats to public safety, and for suits based on over $1 billion in tangible damages. The criminal liability is aimed at punishing companies who lie on their permit applications or who knowingly violate the conditions of their permits. The goal is to deter AI developers who recklessly release profoundly dangerous technology without inspiring frivolous litigation.

Section 18 provides for the orderly activation and use of emergency powers in case the government becomes aware of a clear and present danger based on highly advanced AI. The government would be able to shut down and sequester AI systems that were actively threatening the public, with appropriate compensation paid to innocent bystanders. ***Section 19*** offers protections for whistleblowers, including whistleblowers who call attention to the need for the government to use these emergency powers.

Sections 20 through 24 conclude by tackling some of the bill's administrative details, like inter-agency cooperation, preemption, funding authorization, and severability.

Sections 1 through 3 – Short Title, Sense of Congress, and Definitions

- **Section 1**, the Short Title, gives other lawyers a convenient way of citing our bill in future laws. It also includes a Table of Contents that lists all the sections in the bill.
- **Section 2**, the Sense of Congress, is a political statement that explains why the bill is necessary and what the bill is trying to accomplish.
- **Section 3**, the Definitions section, defines all the special terms used by the bill in alphabetical order.
    - The bill is designed to protect against "major security risks" posed by "frontier AI."
    - A "major security risk" is defined as including any of:
        - the existential risks or global catastrophic risks defined in the 2023 NDAA;
        - a serious threat to critical infrastructure, national security, or public safety; or
        - risks that AI systems will establish self-replicating autonomous agents or otherwise evade or defeat human control.
    - Frontier AI is either a major hardware cluster with over 3,000 specialized AI chips, or extremely advanced AI software that:
        - is a general-purpose model developed by a company controlling at least 10^{20} FLOP/s that used at least 10^{26} FLOP on that model;
        - is a general-purpose model developed by a company controlling at least 10^{18} FLOP/s that used at least 10^{26} FLOP on that model and that failed a set of automated benchmarks designed to detect dangerous capabilities;
        - is one of the 10 most powerful AI systems in the world based on a set of benchmarks selected and published by FAIA; or
        - has been found by FAIA's administrator to pose major security risks that are not otherwise widespread in the AI industry, based on substantial evidence published in the Federal Register.
    - **Flexible and adjustable criteria**: Systems that technically qualify as frontier AI but that clearly pose minimal risks will be exempted from the permit requirement based on a "fast track" exemption form (see Section 11). Conversely, AIs that ought to qualify as frontier AI but fall short of the formal criteria can be included in the bill's permit scheme over time as the agency updates the definition of frontier AI to keep pace with changes in technology.

Section 4 – Creating a New Administration

- **Responsibilities**: The bill would create a new federal agency, called the "Frontier Artificial Intelligence Administration."

- **No Cabinet Department**: As written, the bill does not place this agency under any Cabinet department, nor is the Administrator a cabinet-level position. The Administration is independent and directly accountable to the President, like the NSF, CIA, or NASA.
- **The Administrator:** The President is encouraged to nominate a leader with demonstrated experience in securing advanced technology, e.g., biosecurity or cybersecurity. The leader must be confirmed by the Senate.
- **Hiring Bonuses**: It can be difficult to recruit technical staff on a government salary, so the bill provides for a 50% pay increase over the default civil service salaries, and gives the Administrator flexible hiring authority to make direct hires without going through the lengthy and cumbersome civil service hiring process.
- **Conflicts of Interest:** Senior managers can not work for the Administration if they are on an AI company's board, own stock in an AI company (other than via a mutual fund or blind trust), have worked for an AI company within the last 3 years, or have an immediate family member with such a conflict. The Administrator can waive the last two requirements for other officers, but the waiver gets published in the Federal Register.
- **Volunteers and Donations:** the bill explicitly allows volunteers to work for the Administration and donors to donate to the Administration, as long as there is no real or apparent conflict of interest.

Section 5 – Deputy Administrators

- **The Deputy Administrators:** The Administrator is required to appoint Deputy Administrators to oversee each of FAIA's divisions. The Administrator can remove these Deputies, but must promptly appoint a replacement and publish the cause (or lack thereof) in the Federal Register.
- **Divisions:** Each Deputy Administrator would take charge of one of the core functions of the office, including:
    - *Monitoring*: Monitoring concentrations of advanced AI hardware and tracking down any suspicious or missing information about who is using that hardware.
    - *Standards*: Designing and implementing permit requirements that check whether specific AI systems' deployment would pose major security risks.
    - *Enforcement*: Enforcing the law's civil and criminal liability provisions against those who develop or deploy dangerous AI systems.
    - *Grants Management*: Awarding and evaluating grants to support public access to compute, research into hardware safety features, development of improved safety evaluations, and voluntary internal audits and red-teaming for small businesses.
    - *Public Interest*: Consulting with non-governmental organizations, investigating and reporting on frontier AI systems of special concern and playing 'devil's advocate' to argue that some of the riskiest systems should have their permits denied.

Section 6 – Updating the Thresholds for Frontier AI

- **Rulemaking Authority:** Like other federal agencies, the Frontier AI Administration will be able to publish its own regulations that flesh out the details of how the Responsible AI Act gets implemented.
- **Updating Definitions:** The bill specifically gives the Administration the power to update all technical definitions in the bill, including the definition of frontier AI. The Administration is encouraged to define technical thresholds for biochemical data and other types of specialized data that could be dangerous out of proportion to their size.
- **Public Participation:** if the Administration is behind schedule on completing one of its responsibilities, or if a citizen has an idea for a new rule that would promote AI safety, then anyone can petition the Administration to take action. The Administration must offer a written response to these petitions within 60 days; the responses are subject to judicial review.
- **Major Questions Doctrine:** the major questions doctrine has recently been used to cast doubt on federal agencies' ability to take on ambitious new projects. Part of this section makes it clear that the Administration really does have Congress's permission and authority to carry out all of its responsibilities.

Section 7 – Reporting Significant AI Hardware Transactions

- **Self-Reporting Requirement:** The bill includes a requirement for anyone who buys, sells, destroys, transports, or otherwise transfers at least 100 "high-performance AI chips" in the same quarter to report their transactions using a website that the Administration will set up.
    - This represents about $3 million worth of equipment.
    - The relatively short form will ask for basic information like the address at which the chips are stored and the general purpose for which they are being used.
    - This will help the government keep track of the stock and flows of these strategic resources and give the US a chance to detect any important concentrations of AI chips that have unknown owners or unknown purposes.
- **Definition of High-Performance Hardware**: This is currently defined based on the processing performance and density standards set by the 10/25/23 Advanced Computing Chips Rule (AC/S IFR); it includes the A100 and H100. The Administration can update these standards over time.

Section 8 – AI Hardware Security Audits

- **Major AI Hardware Clusters:** The bill includes a requirement for anyone who owns or operates a "major AI hardware cluster" to pass an annual security audit. This applies only

to data centers with at least 3,000 high-performance chips, i.e., more than about $90 million in equipment, not counting the cost of the building and infrastructure.

- **What the Audit Covers:** The audit will check to make sure that the hardware cluster meets minimum cybersecurity requirements and physical security requirements (to prevent the chips from being hacked or stolen), and that the cluster's owners can reliably identify the real identities of the cluster's major customers.
- **How Audit Evaluated:** The audit report must be sent to FAIA, which will have a chance to require additional safeguards or reject the permit if the audit report is incomplete, inadequate, or unable to vouch for the data center's security.

Section 9 – Automated Benchmarks for Medium-Size Developers

- **Medium-Compute AI Developers:** AI developers who controlled between 10^{18} and 10^{20} FLOP/s during the last quarter and who used that compute to complete a training run of a large, advanced, general-purpose AI model must check that model against a set of automated benchmarks before publicly deploying it. At current market prices, this would only apply to developers who own at least $50 million in advanced chips, or who are spending at least $20 million per year on renting those chips.
- **Automated Benchmarks:** The benchmarks would be designed to check for dangerous capabilities, and they must be privacy-preserving and substantially automatic. Medium-size developers will not be asked to share their model weights with the government, nor will they be asked to wait for more than 24 hours while the benchmarks are evaluated. If there are no benchmarks available that meet these criteria, then the Administration will offer grant funding to try to develop such benchmarks, and medium-size developers will be exempt from the requirement until adequate benchmarks are available.
- **Discovery of Dangerous Capabilities:** If the benchmarks show that an AI system created by a medium-size developer poses major security risks, then the developer must pause or cancel commercial access to the system to the extent technically feasible, and apply for and receive a permit based on an AI software security audit before they continue developing any similar system.

Section 10 – AI Software Security Audits

- **Frontier AI Systems:** Software audits are only required for frontier AI systems, which usually means large, general-purpose, advanced AI systems trained by developers who controlled at least 10^{20} FLOP/s of compute – about $5 billion worth of chips.
- **Contents of Audit:** The audit report must include explicit statements about:
    - the auditor's financial independence and technical qualifications,
    - the length of time and the degree of access that the auditor had for the model being tested,

    - the auditor's opinion about what major security risks might be posed by the model and the degree to which each of these risks has been adequately mitigated.
- **Timing of Requirement:** AI systems that have already been deployed are not required to undergo an audit. Otherwise, an audit is needed before the system's first "significant deployment," defined as using the AI system to generate more than 1 gigabyte or 100 million tokens of output (approximately the amount of text on a large wall of full bookshelves) for any purposes other than safety research and internal evaluation.
- **How Audit Evaluated:** The audit report must be sent to FAIA, which will have a chance to require additional safeguards or reject the permit if the audit report is incomplete, inadequate, or unable to show that deploying the AI system will be very unlikely to pose major security risks.

Section 11 – Permit Application Forms

- **Most AI Models Are Exempt:** Most AI developers won't need to fill out any forms at all, because their models won't meet the definition of frontier AI, and so they will continue to be totally unregulated.
- **Fast-Track Exemptions:** AI developers with models that technically qualify as frontier AI but that obviously don't pose any major security risks still won't have to go through the full audit and permitting process. Instead, FAIA is ordered to design a two-page form that will let AI tools like self-driving cars, fraud detection systems, and recommender engines carry on with their work, even if those models are so large that they might be swept up in the current definition of frontier AI.
- **Initial and Renewal Applications:** FAIA must promptly develop and publish application forms for initial applications and renewal applications for frontier AI permits. Some questions will be different for hardware permits vs. software permits.
- **Application Fees:** FAIA may charge application fees to cover the cost of the permitting process. However, researchers, open source developers, and anyone applying for a fast-track exemption must be exempt from all such fees. If fees are charged, a percentage of them must be set aside to provide assistance to small businesses to help them complete their applications.

Section 12 – Recommended Scoring Factors

- **Suggested Rubrics:** The legislation includes suggested rubrics for auditors to use when evaluating a frontier AI system. Auditors are not required to use all of these factors – there may be a good reason why one of them does not apply to a particular case – but if several factors are inexplicably missing, then that could contribute to a finding by FAIA that the audit was incomplete, arbitrary, or capricious.
- **Hardware Scoring Factors:** The suggested scoring factors for frontier AI hardware audits include the data center's physical security plan, its cybersecurity plan, its know

your customer (KYC) protocol, and its plan for setting aside a portion of its compute to offer for sale to small businesses and entrepreneurs.

- **Software Scoring Factors:** The recommended scoring factors for frontier AI software include:
    - the developer's ability to accurately forecast the capabilities of its model at varying levels of resources,
    - the developer's plan for constructively responding to unexpectedly advanced capabilities,
    - the developer's plan for safety testing and research (including the fraction of compute made available for such a plan),
    - the developer's plan for managing the security of its model weights and other sensitive data, including by controlling who has access to each type of data and how unauthorized copying will be detected and prevented,
    - the extent to which the software might be able to autonomously survive and spread,
    - the extent to which the software might contribute to biological, nuclear, or cyber-offense risks,
    - the developer's liability insurance, and
    - the developer's track record and reputation for accurately predicting and reporting on the capabilities, risks, and safeguards associated with its AI systems.
- **Accommodates Open-Source and Closed-Source Models:** The recommended scoring factors make room for open-source, closed-source, and hybrid designs. Instead of favoring one paradigm over the other, the scoring factors ask how the developer plans to prevent major security risks from arising using their preferred paradigm. If the developer plans to publish their model weights, then they must explain why these weights will not pose major security risks even after they are subject to hostile fine-tuning. If the developer plans to keep their model weights private, then they must explain how they will do so, and how they will respond if the weights are nevertheless leaked.

Section 13 – Adjudicating Permit Applications

- **AI Permit Judges:** After a company submits a permit application, it will be evaluated by a panel of 2 AI Permit Judges (AIPJs) and the Deputy Administrator for Public Interest. The AIPJs will have both scientific and legal knowledge, although they will not necessarily be lawyers. They will work directly for the Administration.
- **Recommendation of Approval by Default:** If the AIPJs and the Deputy Administrator do not object within 60 days after receiving an application, then it is automatically recommended to the Administrator for final approval. Otherwise, the evaluators will meet and confer and attempt to agree on a recommendation, which could include unconditional approval, approval with conditions, an order to revise and resubmit the application, or a

rejection. If the evaluators cannot agree, then the application is recommended for rejection.

- **Administrator Makes Final Decision:** After reviewing and considering the recommendation of the evaluators and their written opinion, the Administrator makes a final decision about how to rule on each application.
- **Tight Timelines:** All the steps of the adjudication and appeals process are designed with short, firm deadlines so that AI developers can get a final decision quickly enough to move forward with their legitimate business needs.
- **Updating the Standards:** Like the definition of frontier AI, the Administration can update its rubrics and application forms at any time. A permit that a company received under an older standard would still be valid unless the Administration uses its emergency powers (see Section 18).

Section 14 – Appealing Permit Applications

- **Appeals Board:** the applicant, the Administrator, or the Deputy Administrator for Public Interest can appeal a decision of the AIPJs to a special Appeals Board within the Administration, made up of 7 experts with diverse professional skills, including lawyers, scientists, and risk management experts.
- **Administrator's Final Correction:** the Administrator can personally overrule the Appeals Board; if they do, they must explain why the Appeals Board's decision does not further the purposes of the bill and publish the explanation in the Federal Register.
- **Judicial Review by DC Circuit Court:** the Federal Court of Appeals can review the Administration's final word on any given application at the request of an applicant or the Deputy Administrator for Public Interest. The usual appeals standard would apply: the Court can overturn the Administration's decision if it relied on the wrong legal principles or if its application of those principles was arbitrary or capricious.

Section 15 – Analysis and Reporting

- **Monthly Trends Analysis:** Using the data from self-reported high-performance hardware transactions, licensing data, and general research on manufacturing and industry trends, the Deputy Administrator for Monitoring is required to put together a monthly report for the government's internal use, tracking where large concentrations of compute are located, who's using it, what they're using it for, and taking note of anything suspicious.
    - If there's anything that doesn't add up, the Administration has broad subpoena power to go investigate and take any evidence they might need to find out what happened to the missing chips.
    - Not complying with the subpoenas would be a crime.

- **Bulletin in Federal Register:** Each year, the Administration has to publish a bulletin in the Federal Register showing their waivers of conflicts of interest, their salary increases, their use of emergency powers, and similar news items. The idea is that if they're doing anything nefarious, it will get some attention based on these publications, or at least it makes it easier for watchdog groups to keep an eye on them.
- **Report to Congress:** Each year, the Administration also has to send a report to Congress on the Administration's work and on the current state of major security risks from AI. The report is limited to 20 pages. If Congress asks for more information, then of course the Administration will provide it.
- **Quarterly Briefings for OSTP:** A senior official from the Administration will personally brief OSTP four times a year on the latest threats and developments in frontier AI.
- **Other Reports to Government:** The Administration will take steps to keep other agencies informed about major security risks from frontier AI, especially when such risks are about to increase or when the government is about to make an important decision related to these risks.

Section 16 – Civil Liability

- **Duty of Care:** The bill makes it explicit that everyone working on frontier AI owes a duty to everyone in the United States to make sure that the AI does not cause harm to innocent bystanders, does not autonomously spread to other servers without permission, and cannot be easily misused by third parties.
- **Joint and Several Liability:** If multiple defendants all contributed to causing the same kind of harm, then a plaintiff can pick and choose which one(s) to sue and recover the full amount from any or all of them. This is helpful when the person who was most to blame is not the person with the deepest pockets.
- **Tangible damages:** This section provides more aggressive remedies for 'tangible damages,' meaning wrongful death, physical injury or illness, direct financial losses, and damage to physical property. Emotional distress and invasion of privacy are not considered 'tangible damages.'
- **Private Right of Action:** People who suffered at least $1 billion in tangible damages from frontier AI can directly sue the wrongdoers, without having to rely on the government to sue for them.
- **Public Right of Action:** The Administrator can also choose to sue anyone who violates the duty of care in this section, regardless of the amount or kind of damages. The Administrator may attempt to recover damages on behalf of the public at large, or on behalf of identified victims, or may simply seek injunctive relief.
- ***Per se* violations:** If a frontier AI developer fails to get a frontier AI permit, lies on its permit application, or violated the terms of its permit application, and that causes harm, then the developer can be sued for that harm even if the plaintiff can't identify precisely

what was wrong with the defendant's AI. In that case, the defendant is strictly liable for any tangible damages.

- **Civil Penalty:** If found guilty, frontier AI developers and deployers can be fined up to $1 million per day for deliberate ongoing violations of the duty of care, to encourage companies to stop breaking the rules once they get caught.
- **Foreseeability of Harm:** The fact that the specific way that an AI became unreliable was a surprise to its developers is not a valid defense; such developers are still liable because they knew or should have known that frontier AI poses a wide variety of severe risks, some of which may not be detectable in advance.
- ***Ex ante* Punitive Damages:** A plaintiff who is harmed by frontier AI in a way that suggests that the AI narrowly avoided causing a catastrophe can sue for "*ex ante*" punitive damages that are designed to force frontier AI developers to internalize some of the costs imposed on society by their dangerous technologies. If an AI system fails badly enough, then its designer may go bankrupt, or, in extreme cases, the court system might fail altogether. Because AI developers are not expected to be made to pay ***after*** the fact for the full amount of their damages in these extreme cases, they should be made to pay a portion of those damages ***before*** they occur based on documented evidence of near-misses.

Section 17 – Criminal Liability

- **Failure to Self-Report Hardware:** A wholesaler who fails to self-report their high-volume trading in high-performance GPUs commits a criminal **infraction**, like running a red light. Technically they could spend a few days in jail, but unless they're a repeat offender, it would usually get handled by a fine of several hundred dollars.
- **Licensing Misdemeanors:** Misrepresenting the safety of frontier AI, or building frontier AI without a permit, or recklessly violating the terms of a frontier AI permit is a misdemeanor, punishable by several months in jail and $100,000 in fines per person.
- **Licensing Felonies:** Directly lying on a frontier AI application, or intentionally violating the terms of your frontier AI permit, or intentionally refusing to apply for a frontier AI permit after receiving notice that one is required, or failing to comply with an emergency order from the Administrator is a felony, punishable by several years in jail and $250,000 in fines per person.
- **Other Crimes:** Similar crimes have similar punishments, even if they're not about permitting. *Attempting* to commit a frontier AI crime (but not following through) gets a punishment that's one step milder than the penalty for the underlying crime.
- **Corporate Penalties:** corporations that commit a misdemeanor have their permits suspended for 6 months, and corporations that commit a felony have their permits canceled, have to sell or destroy their hardware and model weights, and can't apply for a new permit for 5 years. Corporations aren't allowed to give their staff extra pay to relieve the sting if their staff get fined for personally carrying out illegal acts.

- **Non-Profit Enforcement:** Normally, criminal acts are fined in proportion to the profits received from the crime or the losses inflicted by the crime. For open source or non-profit groups that might not be directly earning money, there's a flat penalty available of up to $2 million for a misdemeanor or $25 million for a felony.
- **Flexible Prosecution:** as with the civil violations, these crimes can be prosecuted directly by the Administration, or the Administration can delegate that work to the Justice Department. The statute of limitations is 2 years for infractions, 5 years for misdemeanors, and 10 years for felonies.

Section 18 – Emergency Powers

- **Triggering an Emergency:** If the President finds that frontier AI is posing a "major security risk", or if the Administrator finds that AI is posing a "clear and imminent major security risk that cannot be reliably prevented through ordinary civil and criminal enforcement," then either of them can declare a state of emergency, which immediately activates a suite of emergency powers.
- **Administrator's Powers:** The Administrator can suspend a frontier AI permit, issue restraining orders, encrypt model weights, require people to take additional safety precautions, and generally impose a moratorium on further AI research and development. These powers last for 60 days, or longer if confirmed by the President.
- **Presidential Powers:** The President has all of the powers of the Administrator, plus they can also destroy AI hardware, delete model weights, permanently cancel permits, and physically seize AI laboratories with guards to directly prevent companies from accessing their own labs. These powers last for 1 year, or longer if confirmed by Congress.
- **Serving Notice:** Emergency powers kick in when people have actual notice of an order, or 72 hours after the order is issued, whichever comes later. This gives the Administrator an incentive to make sure people actually find out about the order, so that the order will be quickly enforceable. Individual labs would get "served" by couriers, and a general moratorium would be announced by television and radio.
- **Compensation:** People are eventually entitled to compensation for any damage caused by an emergency order, but only for direct and actual losses, not for the expected value of their counterfactual profits. You still have to comply with the order while you're waiting to get paid.
- **Judicial Review:** Federal district courts can overturn a declaration of emergency, but only if they find "clear and convincing evidence" that either there are no major security risks, or that the Administrator or President has exceeded their lawful authority.
- **Publication in Federal Register:** After activating its emergency powers, the Agency needs to write a report about what happened in its annual bulletin in the Federal Register.

Section 19 – Whistleblower Protection

- **Who Qualifies:** anyone who speaks out against, reports, or refuses to participate in any practice forbidden by the Act can qualify as a whistleblower.
- **Good Faith Belief:** a whistleblower is still protected even if they're wrong about whether a practice was forbidden by the Act, as long as they had a reasonable belief that the Act was being violated.
- **Comprehensive Protection:** employers cannot fire, demote, harass, or otherwise take any action against a whistleblower based on the whistleblower's reports, with three narrow exceptions:
    - Whistleblowers can be suspended for 1 month with full pay during an investigation.
    - Whistleblowers can be punished for unrelated matters if the employer can demonstrate that there was actual misconduct and that the penalty for that misconduct was typical and reasonable.
    - A company can discharge a whistleblower by giving them 2 years of severance pay and a neutral reference for future employers.
- **Remedies:** punishing a protected whistleblower outside of these exceptions is a crime, and the whistleblower can also bring a civil suit for reinstatement and back pay.

Sections 20 through 24 – Miscellaneous

- **Section 20** directs other agencies to cooperate with the AI Administration, especially around antitrust, which could otherwise allow two companies to merge in ways that increase AI risks.
- **Section 21** states that the bill only preempts weaker state laws – if one of the 50 states wants to pass its own AI safety laws, then that's fine, as long as they're at least as strong as this bill.
- **Section 22** specifies that the Administration is allowed to spend whatever funding is authorized for it by Congress, as well as money it receives from licensing fees, fines, and donations. **Section 23** supports this permission by establishing an "AI Safety and Security Fund" that will hold the money collected until it is lawfully expended.
- **Section 24** is the 'severability' clause – if the Supreme Court strikes down part of this bill as unconstitutional, then this clause instructs the Supreme Court to rescue the rest of the bill using any means necessary, including writing their own replacement section(s).