The Responsible Advanced AI Act:
Section-by-Section

Overview

The Center for AI Policy 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, errors, 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 Advanced Artificial Intelligence Act ("RAAIA") 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 us at info@aipolicy.us if you would like help in writing a standalone bill based on one or two sections of the RAAIA or otherwise adapting this work to meet your office's needs.

Structure of the Bill

Sections 1 through 3 of the RAAIA define "artificial intelligence" and the "major security risks" that AI might soon pose. AI is classified into four tiers of concern based on how likely it is to generate major security risks. AI systems in the lowest tier of concern are so unlikely to pose catastrophic risks that we do not see a need to have the government screen them for threats to public safety, and AI systems in the highest tier of concern have the potential to cause such grave harm that we do not think they should be released unless their developers can affirmatively demonstrate that they will be safe.

Section 4 of the RAAIA sets up a new federal office called the Frontier AI Systems Administration ("FAISA") 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 through 9 explain how AI developers would be required to take steps to safeguard the AIs they are training. AIs that pose only a moderate concern would go through a quick automatic test to see if they are unusually advanced for their size; if not, they would be automatically approved. Developers of AIs that pose a high level of concern would need to explain what safeguards they are taking to protect the public and then apply for a permit to train or deploy their AI based on that safety plan.

Sections 10 and 11 describe how the permit applications will be evaluated, and, if necessary, appealed. The bill sets up panels of "AI law judges" who will have expertise in both law and technology and who can make recommendations based on safety rubrics. To protect American innovation, developers will have a legal right to have their applications heard and ruled on quickly, typically within 90 days after they are submitted, and they will be able to continue operations while an application is pending.

Section 12 tasks the government with tracking and monitoring concentrations of high-powered microchips that have been specialized for use in AI systems. ***Section 13*** then requires the government to use this and other information to develop reports on trends in the use and capabilities of general-purpose AI and to advise other offices on how to respond to these trends.

Sections 14 and 15 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 $100 million in tangible damages. The criminal liability is primarily 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 16 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 17*** offers protections for whistleblowers, including whistleblowers who call attention to the need for the government to use these emergency powers.

Sections 18 through 21 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 from the 2023 NDAA,
        - a serious threat to essential infrastructure, national security, or public safety
        - risks that AI systems will establish self-replicating autonomous agents or otherwise permanently escape from human control
    - Frontier AI is defined as Tier 3 and Tier 4 of a four-tiered classification system.
        - *(Tier 1) Low-concern AI***:** AI trained on less than 10^24 FLOP is not regulated.
        - *(Tier 2) Medium-concern AI***:** AI systems trained on more than 10^24 FLOP but less than 10^26 FLOP are typically considered medium-concern. Anyone developing a medium-concern AI must pre-register its training run, and score the AI on automated performance benchmarks. If the AI achieves an unexpectedly high score on these benchmarks, it can be reclassified as high-concern AI.
        - *(Tier 3) High-concern AI***:** AI systems trained on more than 10^26 FLOP are typically considered high-concern; their designers must submit a safety plan and get a permit before the AI system is trained, deployed, or proliferated.
        - *(Tier 4) Extremely high-concern AI***:** AI systems that exhibit signs of having catastrophic capabilities, such as assisting with the design of WMDs or automatically hacking into millions of computers, are considered extremely high-risk. In addition to passing the standard safety evaluation, such systems must be accompanied by strong evidence that they are fundamentally reliable, such as mathematical proof of alignment.
    - **Flexible and adjustable criteria**: Systems that technically qualify as "frontier AI" but that clearly pose minimal risks can be exempted from the permit requirement based on a "fast track" exemption form (see Section 8). 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 doesn't 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't 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.
- **Paperwork Exemption**: normally, any new form issued by a federal agency that collects any information from the public has to be pre-approved by the Office of Management and Budget (OMB). To speed up the Administration's work, it can issue forms immediately, unless and until OMB tells it to revise those forms.

Section 5 – Deputy Administrators

- **The Deputy Administrators:** The Administrator is required to appoint Deputy Administrators to oversee each of FAISA's divisions. The Administrator can remove these Deputies for cause, but must promptly appoint a replacement and publish the cause 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 are safe enough to be trained or deployed.
    - *Enforcement:* Enforcing the law's civil and criminal liability provisions against people who develop or deploy dangerous AI systems.

- *Algorithms:* Keeping track of algorithmic progress, updating the definition of "frontier AI," and briefing Congress & the White House on progress in advanced AI.
- *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" and the thresholds for the four categories of permits (see Section 8).
- **Varying Deadlines:** Typically, updating a definition takes at least 45 days; the agency must go through a "notice and comment" process that allows anyone with an opinion to weigh in and requires the agency to at least acknowledge those opinions in the Federal Register.
  - If the Administrator offers a very short (500 words or less) change whose only effect is to update the definitions based on changes in available technology, then it can pass with only 10 days' notice.
  - If the Administrator wants to offer binding regulations about best practices in frontier AI (along the lines of the NIST AI RMF), then that requires holding at least 2 public listening sessions with at least 90 days' notice.
- **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.
- **Full Support of Congress:** the major questions doctrine and the Congressional Review Act have 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, including updating the thresholds for frontier AI, and that the courts should not set aside a lawful rule unless they have clear and convincing evidence that it would be safe to do so.

Section 7 – Pre-Registration for Medium-Concern AI

- **How to Pre-Register:** If you're training a medium-concern AI, then once a month or at the conclusion of each training run (whichever comes first), you have to run some automatic performance benchmark tests and log the results on a government website. You would also include your name, your address, the general purpose of the training run, and the total amount of compute being used for the training run.
- **Surprisingly High Performance:** If your medium-risk AI scores an average of more than 80% on the performance benchmarks (which would be unexpectedly high given the 10^26 FLOP cutoff for medium-risk AI) then you have to stop training, begin treating the AI as high-concern, and apply for a permit to train a high-concern AI.

Section 8 – Permit Applications for High-Concern AI

- **Categories of Permits:** The bill sets up four different categories of permits – one for each of owning advanced AI hardware, training frontier AI models, having access to the resulting model weights, and deploying frontier AI models.
    - Hardware needs a permit if it has a combined throughput of 10^{18} FLOP/s.
    - Software needs a permit if it is designed to create frontier AI, or could plausibly create frontier AI, or has created frontier AI, or is itself a frontier AI system.
- **Application Forms:** The Administration is instructed to design and publish an application form for each type of permit; the form is meant to collect enough information to let the Administration make an informed decision about the risks posed by each AI system.
- **Fast-Track Exemption Form:** A fifth type of form, the "fast track" exemption form, gets AI developers who aren't posing any major security risks out from under the bill's authority. The Administration 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 they technically qualify as "frontier AI." AI Systems that qualify for the fast track exemption don't have to participate in the rubrics or judging described in the rest of this section.
- **Grandfather Clause:** AI systems that were already operational before the bill is passed can continue without a permit. AI systems that were in the process of being trained when the bill is passed get 60 days to apply for a permit, and they can continue operations while they wait to see if their permit application is approved.
- **Rubrics for Scoring Applications:** The Administration has to design and publish a rubric that assigns points to each application based on many different factors that are meant to evaluate the risks posed by a given AI system and the likely effectiveness of the precautions that the AI system's owners are taking. Some factors are common to all four types of applications, and some factors are specific to one particular application type. For

a complete list of these factors, please see the bill itself. As an illustration, the factors include:
    - Traceability or watermarking
    - Alignment
    - Kill switches or other means of terminating an unsafe AI
    - Cybersecurity
    - Incident reporting
- **Mandatory Insurance:** All frontier AI developers are required to place and maintain liability insurance to get a permit. The details of what coverage is required will be set by the Administration.
- **Open Source Considerations:** The Administration cannot find that an AI system is inherently safe or inherently dangerous simply because it is open source. Instead, the Administration must weigh the overall effect of approving the training or deployment of each system and determine what additional precautions (if any) are needed to protect the public from the catastrophic risks of that system. Open source projects and projects primarily aimed at conducting research on safety, fairness, or reliability cannot be charged an application fee for their permits.
- **Scheduled Renewals:** Applications will need to be renewed once each year. Training permits will also need to be renewed if there is a new training run that wasn't included in the original permit. Applicants whose renewal forms are still being processed by the Administration can carry on with their old AI development activities, but cannot begin new AI development activities.

Section 9 – Additional Requirements for Extremely High-Concern AI

- **Standards for Identifying Extreme Concern:** The Administration would have 12 months to develop detailed standards for determining whether an AI system is extremely high-concern, based on the system's ability to assist with the development of WMDs, autonomously spread to new servers, destabilize the global balance of power, or otherwise pose catastrophic or existential risks
- **Standards for Verifying Safety of Extremely High-Concern Systems:** The Administration would have 30 months to develop detailed standards for determining whether an AI system is safe enough to be trained despite its theoretical potential for causing extremely severe harms. Such standards could include:
    - Evidence that the AI's architecture is fundamentally safe
    - Mathematical proof that the AI is robustly aligned
    - A demonstration that the AI is inherently unable to contribute to particular types of risks, such as the development of biological or nuclear weapons
- **Affirmative Burden to Demonstrate Safety:** An extremely high-concern AI system cannot be classified as safe enough to train solely because no one has proven that the system is dangerous. Instead, the designers of the system must produce evidence that

conclusively rules out any significant possibility that the AI system could cause a catastrophe.

Section 10 – Adjudicating Permit Applications

- **AI Licensing Judges:** After a company submits a permit application, it will be evaluated by a panel of 3 AI License Judges (AILJs). The AILJs will have both scientific and legal knowledge, although they will not necessarily be lawyers. They will work directly for the Administration.
- **Approval Method:** Each AILJ can only vote to approve an application if it both passes the scoring rubric and seems reasonably safe to the AILJ. A majority of 2 out of 3 AILJs can vote to approve the application with or without conditions, return the application for revisions, or reject it. If there is no majority, the application is rejected.
- **Tight Deadlines:** the process is designed to get applicants through the system as quickly as possible; a permit will typically be approved (if at all) within 90 days of when it is submitted. If the administration is running behind schedule, an applicant will have the right to sue the administration to get a final decision.
- **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 until the end of the year unless the Administration uses its emergency powers (see Section 16).

Section 11 – Appealing Permit Applications

- **Appeals Board:** the applicant, the Administrator, or the Deputy Administrator for Public Interest can appeal a decision of the AILJs 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 Director. The Court might rule that the Administrator's explanation is too arbitrary to stand, or that the Administrator was acting without any legal authority. The bill prevents the Court from approving an application that the Administrator rejected unless there is clear and convincing evidence that this will be safe or that there is no legal basis for the rejection.

Section 12 – Hardware Monitoring

- **Self-Reporting Requirement:** The bill includes a requirement for anyone who buys, sells, destroys, transports, or otherwise transfers any "high-performance hardware" to report the transaction using a website that the Administration will set up.
    - The new owner has to report the transactions within 10 days, or within 24 hours of when they first use the high-performance chip, whichever comes first.
    - The old owner also has to report the transaction within 10 days, giving the Administration a chance to cross-reference transactions and find discrepancies.
- **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). In practice, this means the advanced AI chips such as the A100 and H100, which retail for approximately $30,000 each. The Administration can update these standards over time.

Section 13 – 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 the compute is 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 to help keep Congress interested. 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 14 – Civil Liability

- **Duty of Care:** The bill makes it explicit that everyone working on *any* AI owes a duty to everyone else to make sure that the AI does not cause harm to innocent bystanders. In addition, everyone working on *high-concern* AI owes a duty to everyone else to make sure that the high-concern AI cannot spread to other servers without permission, and that high-concern AI cannot be easily misused by third parties.
- **Joint and Several Liability:** If multiple defendants all contributed to causing you the same kind of harm, you 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.
- **Strict liability:** If your AI causes at least $100 million in physical injury, property damage, and financial losses, then the people you harmed can sue you even if they can't identify precisely what was wrong with your AI.
- **Negligence *per se*:** If you fail to get a license or violate your license terms and that causes harm, then you can be sued for that harm even if the plaintiff can't identify precisely what was wrong with your AI.
- **Civil Penalty:** You can be fined up to $100,000 per day for ongoing violations of the duty of care, to encourage companies to stop breaking the rules once they get caught.
- **Private Right of Action**: People harmed by AI can directly sue the wrongdoers, without having to rely on the government to sue for them, as long as the AI caused at least $100 million in damages.
- **Agency Right of Action:** Conversely, the agency itself can sue a company that is making unsafe AI (or ask the Justice Department to sue that company), even if no private plaintiff wants to step forward. The agency is not subject to the $100 million minimum.
- **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 15 – Criminal Liability

- **Failure to Self-Report Hardware:** Failing to self-report your high-performance GPUs is a criminal **infraction**, like running a red light. Technically you could spend a few days in jail, but unless you're a repeat offender, it usually gets handled by a fine of several hundred dollars.
- **Licensing Misdemeanors:** Misrepresenting the safety of your high-concern AI, or building high-concern AI without a permit, or recklessly violating the terms of a high-concern AI permit is a misdemeanor, punishable by several months in jail and $100,000 in fines per person.
- **Licensing Felonies:** Directly lying on a high-concern AI application, or intentionally violating the terms of your high-concern AI permit, 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 licensing. *Attempting* to commit a high-concern AI crime (but not following through) gets a punishment that's one step milder than the penalty for the underlying crime.
- **Enhanced Corporate Penalties:** corporations that commit these crimes face significantly larger fines, including a penalty equal to twice the money they gained or twice the value of the harm they caused.
    - In case of 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.
    - In addition, 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.
- **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 16 – 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," 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 17 – 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, good faith 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 18 through 21 – Conclusion

- **Section 18** 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 19** 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 20** 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 21** 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).