

GENERAL ASSEMBLY OF NORTH CAROLINA
SESSION 2025

FILED SENATE
Mar 25, 2025
S.B. 735
PRINCIPAL CLERK

S

D

SENATE BILL DRS35250-MCa-139A

Short Title: AI Innovation Trust Fund. (Public)

Sponsors: Senators Salvador, Garrett, and Murdock (Primary Sponsors).

Referred to:

A BILL TO BE ENTITLED
AN ACT TO ENACT THE ARTIFICIAL INTELLIGENCE INNOVATION TRUST FUND.

Whereas, recognizing the rapidly evolving nature of artificial intelligence and the importance of responsible innovation, the General Assembly intends this Act to establish an exploratory, iterative approach to AI governance, inviting stakeholder input and encouraging collaborative development of appropriate and proportionate AI regulations; Now, therefore, The General Assembly of North Carolina enacts:

SECTION 1. Article 10 of Chapter 143B of the General Statutes is amended by adding a new Part to read:

"Part 18A. Artificial Intelligence Innovation.

"§ 143B-472.83A. Artificial Intelligence Innovation Trust Fund.

(a) Fund. – There is established a special, nonreverting fund to be known as the North Carolina Artificial Intelligence Innovation Trust Fund. The Secretary of Commerce shall be the trustee of the fund and shall expend money from the fund to (i) provide grants or other financial assistance to companies developing or deploying artificial intelligence models in key industry sectors or (ii) establish or promote artificial intelligence entrepreneurship programs, which may include partnerships with research institutions in the State or other entrepreneur support organizations. The fund shall consist of appropriations to the Department of Commerce to be allocated to the fund, interest earned on money in the fund, and any other grants, premiums, gifts, reimbursements or other contributions received by the State from any source for or in support of the purposes described in this subsection. Funds in the fund are hereby appropriated to the Department for the purposes set forth in this section, and, except as otherwise expressly provided, the provisions of this section apply to persons receiving a grant or assistance from the fund. Funds provided under this Part shall not support projects involving artificial intelligence intended for mass surveillance infringing constitutional rights, unlawful social scoring, discriminatory profiling based on protected characteristics, or generating deceptive digital content intended for fraudulent or electoral interference purposes.

(b) Definitions. – The following definitions apply in this section:

(1) Advanced persistent threat. – An adversary with sophisticated levels of expertise and significant resources that allow it, through the use of multiple different attack vectors including, but not limited to, cyber, physical or deception, to generate opportunities to achieve objectives including, but not limited to, (i) establishing or extending its presence within the information technology infrastructure of an organization for the purpose of exfiltrating information; (ii) undermining or impeding critical aspects of a mission,



* D R S 3 5 2 5 0 - M C A - 1 3 9 A *

- 1 program or organization; or (iii) placing itself in a position to do so in the
2 future.
- 3 (2) Artificial intelligence. – An engineered or machine-based system that varies
4 in its level of autonomy and which may, for explicit or implicit objectives,
5 infer from the input it receives how to generate outputs that may influence
6 physical or virtual environments.
- 7 (3) Artificial intelligence safety incident. – An incident that demonstrably
8 increases the risk of a critical harm occurring by means of any of the
9 following:
- 10 a. A covered model or covered model derivative autonomously engaging
11 in behavior other than at the request of a user.
- 12 b. Theft, misappropriation, malicious use, inadvertent release,
13 unauthorized access or escape of the model weights of a covered
14 model or covered model derivative.
- 15 c. The critical failure of technical or administrative controls, including
16 controls limiting the ability to modify a covered model or covered
17 model derivative.
- 18 d. Unauthorized use of a covered model or covered model derivative to
19 cause or materially enable critical harm.
- 20 (4) Computing cluster. – A set of machines transitively connected by data center
21 networking of over 100 gigabits per second that has a theoretical maximum
22 computing capacity of at least 10 to the power of 20 integer or floating-point
23 operations per second and can be used for training artificial intelligence.
- 24 (4a) Covered entity. – The legally responsible organization, corporation, or entity
25 that directly oversees and controls the development, deployment, and ongoing
26 operations of a covered model or covered model derivative, including
27 responsibility for compliance with obligations under this Part
- 28 (5) Covered model. – An artificial intelligence model that, due to its scale,
29 application domain, or potential impact, is identified by the Secretary as
30 warranting proportionate regulatory oversight. Factors considered may
31 include, but are not limited to, computing power utilized, model training cost,
32 anticipated scope of application, and foreseeable risks to public safety or
33 individual rights. The Secretary may establish multiple tiers of covered
34 models with corresponding compliance frameworks scaled proportionately to
35 identified risk levels.
- 36 (6) Covered model derivative. – A copy of a covered model that: (i) is
37 unmodified; (ii) has been subjected to post-training modifications related to
38 fine-tuning; (iii) has been fine-tuned using a quantity of computing power not
39 exceeding 3 times 10 to the power of 25 or floating point operations, the cost
40 of which, as reasonably assessed by the developer, exceeds \$10,000,000 if
41 calculated using the average market price of cloud compute at the start of
42 fine-tuning; or (iv) has been combined with other software.
- 43 (7) Critical harm. – A harm caused or materially enabled by a covered model or
44 covered model derivative including: (i) the creation or use in a manner that
45 results in mass casualties of a chemical, biological, radiological or nuclear
46 weapon; (ii) mass casualties or at least \$500,000,000 of damage resulting from
47 cyberattacks on critical infrastructure by a model conducting, or providing
48 precise instructions for conducting, a cyberattack or series of cyberattacks on
49 critical infrastructure; (iii) mass casualties or at least \$500,000,000 of damage
50 resulting from an artificial intelligence model engaging in conduct that acts
51 with limited human oversight, intervention or supervision and results in death,

great bodily injury, property damage or property loss, and would, if committed by a human, constitute a crime specified in any general or special law that requires intent, recklessness or gross negligence, or the solicitation or aiding and abetting of such a crime; or (iv) other grave harms to public safety that are of comparable severity to the harms described herein as determined by the attorney general.

The term does not include harms caused or materially enabled by information that a covered model or covered model derivative outputs if the information is otherwise reasonably publicly accessible by an ordinary person from sources other than a covered model or covered model derivative; (ii) harms caused or materially enabled by a covered model combined with other software, including other models, if the covered model did not materially contribute to the other software's ability to cause or materially enable the harm; or (iii) harms that are not caused or materially enabled by the developer's creation, storage, use or release of a covered model or covered model derivative; provided further, that monetary harm thresholds established pursuant to this section shall be adjusted for inflation annually, not later than January 31, by the growth rate of the inflation index over the preceding 12 months; and provided further, that the inflation index shall consist of the per cent change in inflation as measured by the per cent change in the consumer price index for all urban consumers for the Raleigh metropolitan area as determined by the bureau of labor statistics of the United States Department of Labor.

(8) Critical infrastructure. – Assets, systems and networks, whether physical or virtual, the incapacitation or destruction of which would have a debilitating effect on physical security, economic security, public health or safety in the State.

(8a) Department. – The Department of Commerce.

(9) Developer. – A person that performs the initial training of a covered model by: (i) training a model using a sufficient quantity of computing power and cost; or (ii) fine-tuning an existing covered model or covered model derivative using a quantity of computing power and cost sufficient to qualify as a covered model.

(10) Fine-tuning. – Adjusting the model weights of a trained covered model or covered model derivative by exposing such model to additional data.

(11) Full shutdown. – The cessation of operation of: (i) the training of a covered model; (ii) a covered model controlled by a developer; and (iii) all covered model derivatives controlled by a developer.

(11a) Fund. – The Artificial Intelligence Innovation Trust Fund, as established in this section.

(12) Model weight. – A numerical parameter in an artificial intelligence model that is adjusted through training and that helps determine how inputs are transformed into outputs.

(13) Person. – An individual, proprietorship, firm, partnership, joint venture, syndicate, business trust, company, corporation, limited liability company, association, committee or any other nongovernmental organization or group of persons acting in concert.

(14) Post-training modification. – Modifying the capabilities of a covered model or covered model derivative by any means including, but not limited to, fine-tuning, providing such model with access to tools or data, removing

- 1 safeguards against hazardous misuse or misbehavior of such model or
2 combining such model with, or integrating such model into, other software.
3 (15) Safety and security protocol. – Documented, technical, and organizational
4 protocols that: (i) are used to manage the risks of developing and operating
5 covered models or covered model derivatives across their life cycle, including
6 risks posed by causing or enabling or potentially causing or enabling the
7 creation of covered model derivatives; and (ii) specify that compliance with
8 such protocols is required in order to train, operate, possess or provide external
9 access to the developer's covered model or covered model derivatives.
10 (16) Secretary. – The Secretary of Commerce.
11 (c) Oversight. – The Secretary may convene an AI Innovation and Safety Advisory Panel
12 composed of representatives from industry, academia, civil liberties and consumer advocacy
13 groups, and relevant state agencies. This Panel may provide recommendations, best practices,
14 and advice regarding AI technologies, compliance proportionality, and ethical AI-human
15 collaboration. Recommendations of this Panel shall be publicly accessible and may inform future
16 regulatory proposals.
17 (d) Standards. – The Secretary may consider relevant provisions, guidelines, frameworks,
18 and standards established by the U.S. National Institute of Standards and Technology (NIST),
19 and comparable frameworks, such as the EU AI Act, when developing proposals and
20 recommendations pursuant to this Part.
21 **"§ 143B-472.83B. Requirements for developers of covered models.**
22 (a) Reserved.
23 (b) Reserved.
24 (c) Before beginning to train a covered model, a developer shall do all of the following:
25 (1) Implement reasonable administrative, technical and physical cybersecurity
26 protections to prevent unauthorized access to, misuse of or unsafe
27 post-training modifications of the covered model and all covered model
28 derivatives controlled by the developer that are appropriate in light of the risks
29 associated with the covered model, including from advanced persistent threats
30 or other sophisticated actors.
31 (2) Implement the capability to promptly enact a full shutdown.
32 (3) Implement a written and separate safety and security protocol that: (i)
33 specifies protections and procedures that, if successfully implemented, would
34 comply with the developer's duty to take reasonable care to avoid producing
35 a covered model or covered model derivative that poses an unreasonable risk
36 of causing or materially enabling a critical harm; (ii) states compliance
37 requirements in an objective manner and with sufficient detail and specificity
38 to allow the developer or a third party to readily ascertain whether the
39 requirements of the safety and security protocol have been followed; (iii)
40 identifies a testing procedure which takes safeguards into account as
41 appropriate to reasonably evaluate if a covered model poses a substantial risk
42 of causing or enabling a critical harm and if any covered model derivatives
43 pose a substantial risk of causing or enabling a critical harm; (iv) describes in
44 detail how the testing procedure assesses the risks associated with
45 post-training modifications; (v) describes in detail how the testing procedure
46 addresses the possibility that a covered model or covered model derivative
47 may be used to make post-training modifications or create another covered
48 model in a manner that may cause or materially enable a critical harm; (vi)
49 describes in detail how the developer will fulfill their obligations under this
50 chapter; (vii) describes in detail how the developer intends to implement any
51 safeguards and requirements referenced in this section; (viii) describes in

- 1 detail the conditions under which a developer would enact a full shutdown
2 account for, as appropriate, the risk that a shutdown of the covered model, or
3 particular covered model derivatives, may cause disruptions to critical
4 infrastructure; and (ix) describes in detail the procedure by which the safety
5 and security protocol may be modified.
- 6 (4) Ensure that the safety and security protocol is implemented as written,
7 including by designating senior personnel to be responsible for ensuring
8 compliance by employees and contractors working on a covered model or any
9 covered model derivatives controlled by the developer, monitoring and
10 reporting on implementation.
- 11 (5) Retain an unredacted copy of the safety and security protocol for not less than
12 five years after the covered model is no longer made available for commercial,
13 public or foreseeably public use., including records and dates of any updates
14 or revisions.
- 15 (6) Conduct an annual review of the safety and security protocol to account for
16 any changes to the capabilities of the covered model and industry best
17 practices and, if necessary, make modifications to such policy.
- 18 (7) Conspicuously publish a redacted copy of the safety and security protocol and
19 transmit a copy of said redacted safety and security protocol to the attorney
20 general; provided, however, that (i) a redaction in the safety and security
21 protocol may be made only if the redaction is reasonably necessary to protect
22 public safety, trade secrets, or confidential information pursuant to any
23 general, special, or federal law; (ii) the developer shall grant to the attorney
24 general access to the unredacted safety and security protocol upon request;
25 (iii) a safety and security protocol disclosed to the attorney general shall not
26 be a public record; and (iv) if the safety and security protocol is materially
27 modified, the developer shall conspicuously publish and transmit to the
28 attorney general an updated redacted copy of such protocol within 30 days of
29 the modification.
- 30 (8) Take reasonable care to implement other appropriate measures to prevent
31 covered models and covered model derivatives from posing unreasonable
32 risks of causing or materially enabling critical harms.
- 33 (d) Before using a covered model or covered model derivative for a purpose not
34 exclusively related to the training or reasonable evaluation of the covered model for compliance
35 with State or federal law or before making a covered model or covered model derivative available
36 for commercial, public or foreseeably public use, the developer of a covered model shall do all
37 of the following:
- 38 (1) Assess whether the covered model is reasonably capable of causing or
39 materially enabling a critical harm.
- 40 (2) Record, as and when reasonably possible, and retain for not less than five
41 years after the covered model is no longer made available for commercial,
42 public or foreseeably public use, information on any specific tests and test
43 results used in said assessment which provides sufficient detail for third
44 parties to replicate the testing procedure.
- 45 (3) Take reasonable care to implement appropriate safeguards to prevent the
46 covered model and covered model derivatives from causing or materially
47 enabling a critical harm.
- 48 (4) Take reasonable care to ensure, to the extent reasonably possible, that the
49 covered model's actions and the actions of covered model derivatives, as well
50 as critical harms resulting from their actions, may be accurately and reliably
51 attributed to such model or model derivative.

1 (e) A developer shall not use a covered model or covered model derivative for a purpose
2 not exclusively related to the training or reasonable evaluation of the covered model for
3 compliance with State or federal law or make a covered model or a covered model derivative
4 available for commercial, public or foreseeably public use if there is an unreasonable risk that
5 the covered model or covered model derivative will cause or materially enable a critical harm.

6 (f) A developer of a covered model shall annually reevaluate the procedures, policies,
7 protections, capabilities and safeguards implemented pursuant to this section.

8 (g) A developer of a covered model shall annually retain a third-party that conducts
9 investigations consistent with best practices for investigators to perform an independent
10 investigation of compliance with the requirements of this section.

11 (1) The investigator shall conduct investigations consistent with regulations
12 issued by the Secretary. The investigator shall be granted access to unredacted
13 materials as necessary to comply with the investigator's obligations contained
14 herein. The investigator shall produce an investigation report including, but
15 not limited to: (i) a detailed assessment of the developer's steps to comply with
16 the requirements of this section; (ii) if applicable, any identified instances of
17 noncompliance with the requirements of this section and any
18 recommendations for how the developer can improve its policies and
19 processes for ensuring compliance with the requirements of this section; (iii)
20 a detailed assessment of the developer's internal controls, including
21 designation and empowerment of senior personnel responsible for ensuring
22 compliance by the developer and any employees or contractors thereof; and
23 (iv) the signature of the lead investigator certifying the results contained
24 within the investigation report; and provided further, that the investigator shall
25 not knowingly make a material misrepresentation in said report.

26 (2) Covered entities shall transmit to the Attorney General a confidential copy of
27 any independent investigator's report conducted under this section. An
28 executive summary outlining compliance status and risk mitigation actions
29 shall be made publicly available, with proprietary, sensitive, or
30 security-related information redacted as necessary.

31 (h) A developer of a covered model shall annually, until such time that the covered model
32 and any covered model derivatives controlled by the developer cease to be in or available for
33 commercial or public use, submit to the attorney general a statement of compliance signed by the
34 developer's chief technology officer, or a more senior corporate officer, that shall specify or
35 provide, at a minimum: (i) an assessment of the nature and magnitude of critical harms that the
36 covered model or covered model derivatives may reasonably cause or materially enable and the
37 outcome of the assessment required by this section; (ii) an assessment of the risk that compliance
38 with the safety and security protocol may be insufficient to prevent the covered model or covered
39 model derivatives from causing or materially enabling critical harms; and (iii) a description of
40 the process used by the signing officer to verify compliance with the requirements of this section,
41 including a description of the materials reviewed by the signing officer, a description of testing
42 or other evaluation performed to support the statement and the contact information of any third
43 parties relied upon to validate compliance.

44 A developer shall submit such statement to the attorney general not later than 30 days after
45 using a covered model or covered model derivative for a purpose not exclusively related to the
46 training or reasonable evaluation of the covered model for compliance with State or federal law
47 or making a covered model or covered model derivative available for commercial, public or
48 foreseeably public use; provided, however, that no such initial statement shall be required for a
49 covered model derivative if the developer submitted a compliant initial statement and any
50 applicable annual statements for the covered model from which the covered model derivative is
51 derived.

(i) A developer of a covered model shall report each artificial intelligence safety incident affecting the covered model or any covered model derivatives controlled by the developer to the attorney general within 72 hours of the developer learning of the artificial intelligence safety incident or facts sufficient to establish a reasonable belief that an artificial intelligence safety incident has occurred.

(j) This section shall apply to the development, use or commercial or public release of a covered model or covered model derivative for any use that is not the subject of a contract with a federal government entity, even if that covered model or covered model derivative was developed, trained or used by a federal government entity; provided, however, that this section shall not apply to a product or service to the extent that compliance would strictly conflict with the terms of a contract between a federal government entity and the developer of a covered model.

(k) The Secretary may develop and propose a tiered compliance framework differentiating obligations based on computing scale, intended applications, societal impact, and organizational size. This framework shall be developed through stakeholder consultations and presented to the General Assembly with recommendations for potential adoption.

(l) A developer or covered entity may remain responsible for foreseeable critical harms arising from misuse or unintended use of a covered model or derivative, irrespective of whether such misuse involved fine-tuning. Covered entities may conduct and document pre-deployment risk assessments to identify and reasonably mitigate foreseeable misuse risks.

(m) Covered entities funded under this Act developing AI systems that significantly impact individuals' rights or access to critical services such as employment, housing, education, or financial products may conduct exploratory algorithmic fairness assessments to detect and mitigate potential bias. These assessments may be shared with stakeholders and the Department to inform future policy development.

(n) Covered entities may voluntarily explore methods for disclosing to end-users when they are interacting with an artificial intelligence system, particularly where the nature of interaction is not immediately obvious. Such entities may also explore labeling content generated by funded AI systems where there is potential for it to be mistaken for human-generated content. Findings from these explorations may be reported to the Department to inform future transparency guidelines.

"§ 143B-472.83C. Requirements for computer resource operators training covered models.

(a) A person that operates a computing cluster shall implement written policies and procedures to do all of the following when a customer utilizes computer resources which would be sufficient to train a covered model:

- (1) Obtain the prospective customer's basic identifying information and business purpose for utilizing the computing cluster including, but not limited to: (i) the identity of the prospective customer; (ii) the means and source of payment, including any associated financial institution, credit card number, account number, customer identifier, transaction identifiers or virtual currency wallet or wallet address identifier; and (iii) the email address and telephone number used to verify the prospective customer's identity.
- (2) Assess whether the prospective customer intends to utilize the computing cluster to train a covered model.
- (3) Maintain logs of significant access and administrative actions consistent with commercially reasonable cybersecurity practices.
- (4) Maintain for not less than seven years, and provide to the attorney general upon request, appropriate records of actions taken under this section, including policies and procedures put into effect.
- (5) Implement the capability to promptly enact a full shutdown of any resources being used to train or operate a covered model under the customer's control.

1 If a customer repeatedly utilizes computer resources that would be sufficient to train a
2 covered model, the operator of the computer cluster shall validate said basic identifying
3 information and assess whether such customer intends to utilize the computing cluster to train a
4 covered model prior to each utilization.

5 (b) A person that operates a computing cluster shall consider industry best practices and
6 applicable guidance from the National Institute of Standards and Technology, including the
7 United States Artificial Intelligence Safety Institute, and other reputable standard-setting
8 organizations.

9 (c) In complying with the requirements of this section, a person that operates a computing
10 cluster may impose reasonable requirements on customers to prevent the collection or retention
11 of personal information that the person operating such computing cluster would not otherwise
12 collect or retain, including a requirement that a corporate customer submit corporate contact
13 information rather than information that would identify a specific individual.

14 **"§ 143B-472.83D. Enforcement.**

15 (a) The attorney general shall have the authority to enforce the provisions of this Part.
16 Except as specifically provided in this Part, nothing in this Part shall be construed as creating a
17 new private right of action or serving as the basis for a private right of action that would not
18 otherwise have had a basis under any other law but for the enactment of this Part. This Part
19 neither relieves any party from any duties or obligations imposed nor alters any independent
20 rights that individuals have under State or federal laws, the North Carolina Constitution or the
21 United States Constitution.

22 The attorney general may initiate a civil action in the superior court against an entity in the
23 name of the State or on behalf of individuals for a violation of this chapter. The attorney general
24 may seek:

25 (1) Against a developer of a covered model or covered model derivative for a
26 violation that causes death or bodily harm to another human, harm to property,
27 theft or misappropriation of property, or that constitutes an imminent risk or
28 threat to public safety that occurs on or after January 1, 2026, a civil penalty
29 in an amount not exceeding (i) for a first violation, five percent (5%) of the
30 cost of the quantity of computing power used to train the covered model to be
31 calculated using the average market prices of cloud compute at the time of
32 training or (ii) for any subsequent violation, 15 percent (15%) of the cost of
33 the quantity of computing power used to train the covered model as calculated
34 herein.

35 (2) Against an investigator for a violation of this Part, including an investigator
36 who intentionally or with reckless disregard violates any of such investigator's
37 responsibilities , or for a person that operates a computing cluster in violation
38 of this Part, a civil penalty in an amount not exceeding (i) twenty-five
39 thousand dollars (\$25,000) for a first offense; (ii) fifty thousand dollars
40 (\$50,000) for any subsequent violation; and (iii) five million dollars
41 (\$5,000,000) in the aggregate for related violations.

42 (3) Injunctive or declaratory relief.

43 (4) Such monetary or punitive damages as the court may allow.

44 (5) Attorney's fees and costs.

45 (6) Any other relief that the court deems appropriate.

46 (b) In determining whether a developer exercised reasonable care in the creation, use, or
47 deployment of a covered model or covered model derivative, the attorney general shall consider
48 all of the following:

49 (1) The quality of such developer's safety and security protocol.

50 (2) The extent to which the developer faithfully implemented and followed its
51 safety and security protocol.

(3) Whether, in quality and implementation, the developer's safety and security protocol was comparable to those of developers of models trained using a comparable amount of compute resources.

(4) The quality and rigor of the developer's investigation, documentation, evaluation and management of risks of critical harm posed by its model.

(c) A provision within a contract or agreement that seeks to waive, preclude, or burden the enforcement of liability arising from a violation of this Part, or to shift such liability to any person or entity in exchange for their use or access of, or right to use or access, a developer's product or services, including by means of a contract or adhesion, shall be deemed to be against public policy and void.

Notwithstanding any corporate formalities, the court shall impose joint and several liability on affiliated entities for purposes of effectuating the intent of this section to the maximum extent permitted by law if the court concludes all of the following:

(1) The affiliated entities, in the development of the corporate structure among such affiliated entities, took steps to purposely and unreasonably limit or avoid liability.

(2) As a result of any such steps, the corporate structure of the developer or affiliated entities would frustrate recovery of penalties, damages, or injunctive relief under this section.

(d) Penalties collected pursuant to this section by the attorney general shall be deposited into the General Fund and subject to appropriation.

"§ 143B-472.83E. Cooperation with Attorney General.

(a) For purposes of this section, the following definitions apply:

(1) Contractor or subcontractor. – A firm, corporation, partnership or association and its responsible managing officer, as well as any supervisors, managers or officers found by the attorney general or director to be personally and substantially responsible for the rights and responsibilities of employees under this section.

(2) Employee. – Any person who performs services for wages or salary under a contract of employment, express or implied, for an employer, including:

a. Contractors or subcontractors and unpaid advisors involved with assessing, managing or addressing the risk of critical harm from covered models or covered model derivatives.

b. Corporate officers.

(b) A developer of a covered model or a contractor or subcontractor of the developer shall not:

(1) Prevent an employee from disclosing information to the attorney general or any other public body, including through terms and conditions of employment or seeking to enforce terms and conditions of employment, if the employee has reasonable cause to believe the information indicates that (i) the developer is out of compliance with the requirements of this section or (ii) an artificial intelligence model, including a model that is not a covered model or a covered model derivative, poses an unreasonable risk of causing or materially enabling critical harm, even if the employer is not out of compliance with any State or federal law.

(2) Retaliate against an employee for disclosing such information to the attorney general or any other public body.

(3) Make false or materially misleading statements related to its safety and security protocol in any manner that would constitute an unfair or deceptive trade practice.

(c) An employee harmed by a violation of this section may petition the court for appropriate relief.

(d) The attorney general may publicly release any complaint, or a summary of such complaint, filed pursuant to this section if the attorney general concludes that doing so will serve the public interest; provided, however, that any information that is confidential, qualifies as a trade secret, or is determined by the attorney general to likely pose an unreasonable risk to public safety if disclosed shall be redacted from the complaint prior to disclosure.

(e) A developer shall provide a clear notice to all employees working on covered models and covered model derivatives of their rights and responsibilities under this section, including the rights of employees of contractors and subcontractors to utilize the developer's internal process for making protected disclosures pursuant to subsection (f). A developer is presumed to be in compliance with the requirements of this subsection if the developer:

(1) At all times posts and displays within all workplaces maintained by the developer a notice to all employees of their rights and responsibilities under this section, ensures that all new employees receive equivalent notice and ensures that employees who work remotely periodically receive an equivalent notice; or

(2) At least annually, provides written notice to all employees of their rights and responsibilities under this section and ensures that such notice is received and acknowledged by all of those employees.

(f) A developer shall provide a reasonable internal process through which an employee, contractor, subcontractor or employee of a contractor or subcontractor working on a covered model or covered model derivative may anonymously disclose information to the developer if the employee believes, in good faith, that the developer has violated any provision of this chapter or any other general or special law, has made false or materially misleading statements related to its safety and security protocol or has failed to disclose known risks to employees. The developer shall conduct an investigation related to any information disclosed through such process and provide, at a minimum, a monthly update to the person who made the disclosure regarding the status of the developer's investigation of the disclosure and the actions taken by the developer in response to the disclosure.

Any disclosure and response created pursuant to this subsection shall be maintained for not less than seven years from the date when the disclosure or response is created. Each disclosure and response shall be shared with officers and directors of the developer whose acts or omissions are not implicated by the disclosure or response not less than once per quarter. In the case of a report or disclosure regarding alleged misconduct by a contractor or subcontractor, the developer shall notify the officers and directors of the contractor or subcontractor whose acts or omissions are not implicated by the disclosure or response about the status of their investigation not less than once per quarter.

"§ 143B-472.83. Reporting and regulation.

The Secretary shall file an annual report not later than January 31 with the General Assembly containing: (i) statistical information on the current workforce population in the business of the development of artificial intelligence and in adjacent technology sectors; (ii) any known workforce shortages in the development or deployment of artificial intelligence; (iii) summary information related to the efficacy of existing workforce development programs in artificial intelligence and related sectors, if any; (iv) summary information related to the availability of relevant training programs available in the State, including any known gaps in such programs generally available to members of the public; and (iv) any plans, including recommendations for legislation, if any, to remedy any such known workforce shortages.

The Secretary shall promulgate regulations for the implementation, administration and enforcement of this Part; provided, however, that the Secretary may convene an advisory board for the purposes of: (i) studying the impact of artificial intelligence on the State, including with

respect to its employees, constituents, private business and higher education institutions; (ii) conducting outreach and collecting input from stakeholders and experts; (iii) studying current and emerging capability for critical harms made possible by artificial intelligence developed or deployed in the State; or (iv) advising the Governor and General Assembly on recommended legislation or regulations related to the growth of the artificial intelligence industry and prevention of critical harms.

Not less than annually, the Secretary shall do all of the following:

- (1) Update, by regulation, the initial compute threshold and the fine-tuning compute threshold that an artificial intelligence model shall meet to be considered a covered model, taking into account: (i) the quantity of computing power used to train models that have been identified as being reasonably likely to cause or materially enable a critical harm; (ii) similar thresholds used in federal law, guidance or regulations for the management of artificial intelligence models with reasonable risks of causing or enabling critical harms; and (iii) input from stakeholders, including academics, industry, the open-source community and government entities.
- (2) Update, by regulation, binding investigation requirements applicable to investigations conducted pursuant to this Part to ensure the integrity, independence, efficiency and effectiveness of the investigation process, taking into account: (i) relevant standards or requirements imposed under federal or State law or through self-regulatory or standards-setting bodies; (ii) input from stakeholders, including academic, industry and government entities, including from the open-source community; and (iii) consistency with guidance issued by the National Institute of Standards and Technology, including the United States Artificial Intelligence Safety Institute.
- (3) Issue guidance for preventing unreasonable risks of covered models and covered model derivatives causing or materially enabling critical harms, including, but not limited to, more specific components of, or requirements under, the duties required under this Part. Such guidance shall be consistent with guidance issued by the National Institute of Standards and Technology, including the United States Artificial Intelligence Safety Institute."

SECTION 2. There is appropriated from the General Fund to the Department of Commerce the nonrecurring sum of seven hundred fifty thousand dollars (\$750,000) for the 2025-2026 fiscal year to accomplish the purposes of this act.

SECTION 3. This act becomes effective July 1, 2025.