SENATE DOCKET, NO. 1909 FILED ON: 1/17/2025

SENATE . . . . . No. 37

The Commonwealth of Massachusetts

PRESENTED BY:

Barry R. Finegold

To the Honorable Senate and House of Representatives of the Commonwealth of Massachusetts in General Court assembled:

The undersigned legislators and/or citizens respectfully petition for the adoption of the accompanying bill:

An Act promoting economic development with emerging artificial intelligence models and safety.

PETITION OF:

NAME:	DISTRICT/ADDRESS:
Barry R. Finegold	*Second Essex and Middlesex*

SENATE DOCKET, NO. 1909 FILED ON: 1/17/2025

SENATE . . . . . No. 37

By Mr. Finegold, a petition (accompanied by bill, Senate, No. 37) of Barry R. Finegold for legislation to promote economic development with emerging artificial intelligence models and safety. Advanced Information Technology, the Internet and Cybersecurity.

The Commonwealth of Massachusetts

In the One Hundred and Ninety-Fourth General Court
(2025-2026)

An Act promoting economic development with emerging artificial intelligence models and safety.

Be it enacted by the Senate and House of Representatives in General Court assembled, and by the authority of the same, as follows:

SECTION 1. Chapter 29 of the General Laws is hereby amended by adding the following new section:-

Section 2GGGGGG. Artificial Intelligence Innovation Trust Fund

(a) There shall be established and set up on the books of the commonwealth a separate fund to be known as the Massachusetts Artificial Intelligence Innovation Trust Fund. The secretary of economic development shall be the trustee of the fund and shall, in consultation with the executive director of the Massachusetts Technology Park Corporation established pursuant to chapter 40J, 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 as enumerated in line 7002-8070 of section 2 of chapter 238 of the Acts of 2024; provided, however, that the secretary may seek the commitment of matching or other additional

funds from private sources before making an expenditure from the fund; (ii) establishment or promotion of artificial intelligence entrepreneurship programs, which may include partnerships with research institutions in the commonwealth or other entrepreneur support organizations; or (iii) provide grants or other financial assistance for research in artificial intelligence through or in partnership with the Massachusetts Technology Park Corporation.

(b) There shall be credited to the fund an amount equal to: (i) any appropriations or other money authorized by the general court and specifically designated to be credited to the fund; (ii) interest earned on any money in the fund; and (iii) any other grants, premiums, gifts, reimbursements or other contributions received by the commonwealth from any source for or in support of the purposes described in subsection (a).

(c) Amounts credited to the fund may be expended without further appropriation. For the purpose of accommodating timing discrepancies between the receipt of revenues and related expenditures, the fund may incur expenses, and the comptroller shall certify for payment, amounts not to exceed the most recent revenue estimate as certified by the secretary of elder affairs, as reported in the state accounting system. Any money remaining in the fund at the end of a fiscal year shall not revert to the General Fund and shall be available for expenditure in a subsequent fiscal year.

SECTION 2. The General Laws are hereby amended by inserting after chapter 93L the following new chapter:-

CHAPTER 93M. Artificial Intelligence Models

Section 1. As used in this chapter, the following terms shall have the following meanings unless the context clearly requires otherwise:

"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, program or organization; or (iii) placing itself in a position to do so in the future.

"Artificial intelligence", an engineered or machine-based system that varies in its level of autonomy and which may, for explicit or implicit objectives, infer from the input it receives how to generate outputs that may influence physical or virtual environments

"Artificial intelligence safety incident", an incident that demonstrably increases the risk of a critical harm occurring by means of:

(i) A covered model or covered model derivative autonomously engaging in behavior other than at the request of a user;

(ii) Theft, misappropriation, malicious use, inadvertent release, unauthorized access or escape of the model weights of a covered model or covered model derivative;

(iii) The critical failure of technical or administrative controls, including controls limiting the ability to modify a covered model or covered model derivative; or

(iv) Unauthorized use of a covered model or covered model derivative to cause or materially enable critical harm.

"Computing cluster", a set of machines transitively connected by data center networking of over 100 gigabits per second that has a theoretical maximum computing capacity of at least 1020 integer or floating-point operations per second and can be used for training artificial intelligence.

"Covered model", an artificial intelligence model which is: (i) trained using a quantity of computing power greater than 1026 integer or floating-point operations, the cost of which exceeds $100,000,000 when calculated using the average market prices of cloud compute at the start of training as reasonably assessed by the developer; or (ii) created by fine-tuning a covered model using a quantity of computing power equal to or greater than 3 times 1025 integer or floating-point operations, the cost of which, as reasonably assessed by the developer, exceeds $10,000,000 if calculated using the average market price of cloud compute at the start of fine-tuning; provided, however, that investment 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 Boston metropolitan area as determined by the bureau of labor statistics of the United States department of labor.

"Covered model derivative", a copy of a covered model that: (i) is unmodified; (ii) has been subjected to post-training modifications related to fine-tuning; (iii) has been fine-tuned using a quantity of computing power not exceeding 3 times 1025 or floating point operations, the cost of which, as reasonably assessed by the developer, exceeds $10,000,000 if calculated using the average market price of cloud compute at the start of fine-tuning; or (iv) has been combined with other software.

"Critical harm", a harm caused or materially enabled by a covered model or covered model derivative including: (i) the creation or use in a manner that results in mass casualties of a chemical, biological, radiological or nuclear weapon; (ii) mass casualties or at least $500,000,000 of damage resulting from cyberattacks on critical infrastructure by a model conducting, or providing precise instructions for conducting, a cyberattack or series of cyberattacks on critical infrastructure; (iii) mass casualties or at least $500,000,000 of damage resulting from an artificial intelligence model engaging in conduct that: (A) acts with limited human oversight, intervention or supervision; and (B) 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; provided, however, that "critical harm" shall not include: (i) 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 Boston metropolitan area as determined by the bureau of labor statistics of the United States department of labor.

"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 commonwealth.

"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.

"Fine-tuning", adjusting the model weights of a trained covered model or covered model derivative by exposing such model to additional data.

"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.

"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.

"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.

"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 safeguards against hazardous misuse or misbehavior of such model or combining such model with, or integrating such model into, other software.

"Safety and security protocol", documented technical and organizational protocols that: (i) are used to manage the risks of developing and operating covered models or covered model derivatives across their life cycle, including risks posed by causing or enabling or potentially causing or enabling the creation of covered model derivatives; and (ii) specify that compliance with such protocols is required in order to train, operate, possess or provide external access to the developer’s covered model or covered model derivatives.

"Secretary", the secretary of technology services and security.

Section 2. (a) Before beginning to train a covered model, a developer shall:

(1) implement reasonable administrative, technical and physical cybersecurity protections to prevent unauthorized access to, misuse of or unsafe post-training modifications of the covered model and all covered model derivatives controlled by the developer that are appropriate in light of the risks associated with the covered model, including from advanced persistent threats or other sophisticated actors;

(2) implement the capability to promptly enact a full shutdown;

(3) implement a written and separate safety and security protocol that: (A) specifies protections and procedures that, if successfully implemented, would comply with the

developer’s duty to take reasonable care to avoid producing a covered model or covered model derivative that poses an unreasonable risk of causing or materially enabling a critical harm; (B) states compliance requirements in an objective manner and with sufficient detail and specificity to allow the developer or a third party to readily ascertain whether the requirements of the safety and security protocol have been followed; (C) identifies a testing procedure which takes safeguards into account as appropriate to reasonably evaluate if a covered model poses a substantial risk of causing or enabling a critical harm and if any covered model derivatives pose a substantial risk of causing or enabling a critical harm; (D) describes in detail how the testing procedure assesses the risks associated with post-training modifications; (E) describes in detail how the testing procedure addresses the possibility that a covered model or covered model derivative may be used to make post-training modifications or create another covered model in a manner that may cause or materially enable a critical harm; (F) describes in detail how the developer will fulfill their obligations under this chapter; (G) describes in detail how the developer intends to implement any safeguards and requirements referenced in this section; (H) describes in detail the conditions under which a developer would enact a full shutdown account for, as appropriate, the risk that a shutdown of the covered model, or particular covered model derivatives, may cause disruptions to critical infrastructure; and (I) describes in detail the procedure by which the safety and security protocol may be modified;

(4) ensure that the safety and security protocol is implemented as written, including by designating senior personnel to be responsible for ensuring compliance by employees and contractors working on a covered model or any covered model derivatives controlled by the developer, monitoring and reporting on implementation;

(5) retain an unredacted copy of the safety and security protocol for not less than 5 years after the covered model is no longer made available for commercial, public or foreseeably public use,, including records and dates of any updates or revisions;

(6) conduct an annual review of the safety and security protocol to account for any changes to the capabilities of the covered model and industry best practices and, if necessary, make modifications to such policy;

(7) conspicuously publish a redacted copy of the safety and security protocol and transmit a copy of said redacted safety and security protocol to the attorney general; provided, however, that a redaction in the safety and security protocol may be made only if the redaction is reasonably necessary to protect public safety, trade secrets as defined in section 2 of chapter 93 or confidential information pursuant to any general, special or federal law; provided further, that the developer shall grant to the attorney general access to the unredacted safety and security protocol upon request; provided further, that a safety and security protocol disclosed to the attorney general shall not be a public record for the purposes of chapter 66; and provided further, that if the safety and security protocol is materially modified, the developer shall conspicuously publish and transmit to the attorney general an updated redacted copy of such protocol within 30 days of the modification; and

(8) take reasonable care to implement other appropriate measures to prevent covered models and covered model derivatives from posing unreasonable risks of causing or materially enabling critical harms.

(b) Before using a covered model or covered model derivative for a purpose not exclusively related to the training or reasonable evaluation of the covered model for compliance

with state or federal law or before making a covered model or covered model derivative available for commercial, public or foreseeably public use, the developer of a covered model shall:

(i) assess whether the covered model is reasonably capable of causing or materially enabling a critical harm;

(ii) record, as and when reasonably possible, and retain for not less than 5 years after the covered model is no longer made available for commercial, public or foreseeably public use, information on any specific tests and test results used in said assessment which provides sufficient detail for third parties to replicate the testing procedure;

(iii) take reasonable care to implement appropriate safeguards to prevent the covered model and covered model derivatives from causing or materially enabling a critical harm; and

(iv) take reasonable care to ensure, to the extent reasonably possible, that the covered model’s actions and the actions of covered model derivatives, as well as critical harms resulting from their actions, may be accurately and reliably attributed to such model or model derivative.

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

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

(e)(1) A developer of a covered model shall annually retain a third-party investigator that conducts investigations consistent with best practices for investigators to perform an independent investigation of compliance with the requirements of this section; provided, however, that an investigator shall conduct investigations consistent with regulations issued by the secretary pursuant to section 7. The investigator shall be granted access to unredacted materials as necessary to comply with the investigator’s obligations contained herein. The investigator shall produce an investigation report including, but not limited to: (i) a detailed assessment of the developer’s steps to comply with the requirements of this section; (ii) if applicable, any identified instances of noncompliance with the requirements of this section and any recommendations for how the developer can improve its policies and processes for ensuring compliance with the requirements of this section; (iii) a detailed assessment of the developer’s internal controls, including designation and empowerment of senior personnel responsible for ensuring compliance by the developer and any employees or contractors thereof; and (iv) the signature of the lead investigator certifying the results contained within the investigation report; and provided further, that the investigator shall not knowingly make a material misrepresentation in said report.

(2) The developer shall retain an unredacted copy of the investigation report for not less than 5 years after the covered model is no longer made available for commercial, public or foreseeably public use. The developer shall conspicuously publish a redacted copy of the investigator’s report and transmit to the attorney general a redacted copy of the investigator’s report; provided, however, that a reaction in the investigator’s report

may be made only if the redaction is reasonably necessary to protect public safety, trade secrets as defined in section 2 of chapter 93 or confidential information pursuant to state and federal law; provided further, that the developer shall grant to the attorney general access to the unredacted investigator’s report upon request; and provided further, that an investigator’s report disclosed to the attorney general shall not be a public record for the purposes of chapter 66.

(f)(1) A developer of a covered model shall annually, until such time that the covered model and any covered model derivatives controlled by the developer cease to be in or available for commercial or public use, submit to the attorney general a statement of compliance signed by the chief technology officer, or a more senior corporate officer, that shall specify or provide, at a minimum: (i) an assessment of the nature and magnitude of critical harms that the covered model or covered model derivatives may reasonably cause or materially enable and the outcome of the assessment required by subsection (b); (ii) an assessment of the risk that compliance with the safety and security protocol may be insufficient to prevent the covered model or covered model derivatives from causing or materially enabling critical harms; and (iii) a description of the process used by the signing officer to verify compliance with the requirements of this section, including a description of the materials reviewed by the signing officer, a description of testing or other evaluation performed to support the statement and the contact information of any third parties relied upon to validate compliance.

(2) A developer shall submit such statement to the attorney general not later than 30 days after using a covered model or covered model derivative for a purpose not exclusively related to the training or reasonable evaluation of the covered model for compliance with state or federal law or making a covered model or covered model derivative available for commercial,

public or foreseeably public use; provided, however, that no such initial statement shall be required for a covered model derivative if the developer submitted a compliant initial statement and any applicable annual statements for the covered model from which the covered model derivative is derived.

(g) 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 occured.

(h) 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.

Section 3. (a) (1) A person that operates a computing cluster shall implement written policies and procedures to do all of the following when a customer utilizes compute resources which would be sufficient to train a covered model:

(i) obtain the prospective customer’s basic identifying information and business purpose for utilizing the computing cluster including, but not limited to: (A) the identity of the prospective customer; (B) 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 (C) the email address and telephone number used to verify the prospective customer’s identity;

(ii) assess whether the prospective customer intends to utilize the computing cluster to train a covered model;

(iii) retain any internet protocol addresses used by the customer for access or administration and the date and time of each access or administrative action;

(iv) maintain for not less than 7 years, and provide to the attorney general upon request, appropriate records of actions taken under this section, including policies and procedures put into effect;

(v) 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.

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

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

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

Section 4. (a) (1) The attorney general shall have the authority to enforce the provisions of this chapter. Except as provided in section 5, nothing in this chapter shall be construed as creating a new private right of action or serving as the basis for a private right of action that would not otherwise have had a basis under any other law but for the enactment of this chapter. This chapter neither relieves any party from any duties or obligations imposed nor alters any independent rights that individuals have under chapter 93A, other state or federal laws, the Massachusetts Constitution or the United States Constitution.

(2) The attorney general may initiate a civil action in the superior court against an entity in the name of the commonwealth or as parens patriae on behalf of individuals for a violation of this chapter. The attorney general may seek:

(i) against a developer of a covered model or covered model derivative for a violation that causes death or bodily harm to another human, harm to property, theft or misappropriation of property, or that constitutes an imminent risk or threat to public safety that occurs on or after January 1, 2026, a civil penalty in an amount not exceeding: (A) for a first violation, 5 per cent of the cost of the quantity of computing power used to train the covered model to be calculated using the average market prices of cloud compute at the time of training;;

or (B) for any subsequent violation, 15 percent of the cost of the quantity of computing power used to train the covered model as calculated herein;;

(ii) against a developer of a covered model or covered model derivative a violation of section 5, a civil penalty as outlined in section 185 of chapter 149;

(iii) against an investigator for a violation of subsection (e) of section 2, including an investigator who intentionally or with reckless disregard violates any of such investigator’s responsibilities under said subsection, or for a person that operates a computing cluster in violation of section 3, a civil penalty in an amount not exceeding: (A) $25,000 for a first offense; (B) $50,000 for any subsequent violation; and (C) $5,000,000 in the aggregate for related violations;

(iv) injunctive or declaratory relief;

(v) such monetary or punitive damages as the court may allow;

(vi) attorney’s fees and costs; and

(vii) any other relief that the court deems appropriate.

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

(i) the quality of such developer’s safety and security protocol;

(ii) the extent to which the developer faithfully implemented and followed its safety and security protocol;

(iii) 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;

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

(c) (1) A provision within a contract or agreement that seeks to waive, preclude or burden the enforcement of a liability arising from a violation of this chapter, 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.

(2) 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 that:

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

(ii) 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.

Section 5. (a) For purposes of this section, the following words shall have the following meanings unless the context clearly requires otherwise:

"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 chapter.

"Director", the director of the department of labor standards as established under section 1 of chapter 23.

"Employee", any person who performs services for wages or salary under a contract of employment, express or implied, for an employer, including:

(i) contractors or subcontractors and unpaid advisors involved with assessing, managing or addressing the risk of critical harm from covered models or covered model derivatives; and

(ii) corporate officers.

"Public body" shall have the same meaning as ascribed to it in section 185 of chapter 149.

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

(i) 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:

(A) the developer is out of compliance with the requirements of this chapter; or

(B) 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;

(ii) retaliate against an employee for disclosing such information to the attorney general or any other public body; or

(iii) 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 under chapter 93A.

(c) An employee harmed by a violation of this section may petition the court for appropriate relief as provided in section 185 of chapter 149.

(d) The attorney general or director may publicly release or provide to the governor any complaint, or a summary of such complaint, filed pursuant to this section if the attorney general or director concludes that doing so will serve the public interest; provided, however, that any information that is confidential, otherwise exempt from the provisions of chapter 66, qualifies as a trade secret under sections 42 to 42G, inclusive, of chapter 93 or is determined by the attorney general or director 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:

(i) 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

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

(f) (1) 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.

(2) Any disclosure and response created pursuant to this subsection shall be maintained for not less than 7 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.

(g) This section shall not be construed to limit any rights or obligations of employees under section 185 of chapter 149 or any other state or federal law.

Section 6. The secretary, in consultation with the secretary of economic development and the secretary of labor and workforce development, shall file an annual report not later than January 31 with the joint committee on economic development and emerging technologies, the joint committee on advanced information technology, the internet and cybersecurity and the joint committee on labor and workforce development containing: (i) statistical information on the current workforce population: (A) in the business of the development of artificial intelligence and (B) in adjacent technology sectors as enumerated in line 7002-8070 of section 2 of chapter 238 of the Acts of 2024; (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 commonwealth, 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.

Section 7. The secretary shall promulgate regulations for the implementation, administration and enforcement of this chapter; provided, however, that the secretary may convene an advisory board for the purposes of: (i) studying the impact of artificial intelligence on the commonwealth, 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 commonwealth; or (iv) advising the secretary, governor and general court on recommended legislation or regulations related to the growth of the artificial intelligence industry and prevention of critical harms; provided further, that not less than annually, the secretary shall: (i) 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 as defined in section 1 of chapter 93M, taking into account: (A) 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; (B) 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 (C) input from stakeholders, including academics, industry, the open-source community and government entities; (ii) update, by regulation, binding investigation requirements applicable to investigations conducted pursuant to subsection (e) of section 2 of chapter 93M to ensure the integrity, independence, efficiency and effectiveness of the

investigation process, taking into account: (A) relevant standards or requirements imposed under federal or state law or through self-regulatory or standards-setting bodies; (B) input from stakeholders, including academic, industry and government entities, including from the open-source community; and (C) consistency with guidance issued by the National Institute of Standards and Technology, including the United States Artificial Intelligence Safety Institute; and (iii) 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 chapter 93M; and provided further, that any 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 3. Section 1 of chapter 93M of the General Laws, as inserted by section 2, is hereby amended by striking the definition of "Covered model" and inserting in place thereof the following:-

"Covered model", an artificial intelligence model which is: (i) trained using a quantity of computing power determined by the secretary pursuant to section 7, the cost of which exceeds $100,000,000 when calculated using the average market prices of cloud compute at the start of training as reasonably assessed by the developer; or (ii) created by fine-tuning a covered model using a quantity of computing power that exceeds a threshold determined by the secretary pursuant to section 7, the cost of which, as reasonably assessed by the developer, exceeds $10,000,000 if calculated using the average market price of cloud compute at the start of fine-tuning; provided, however, that investment 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 Boston metropolitan area as determined by the bureau of labor statistics of the United States department of labor.

SECTION 4. Said section 1 of said chapter 93M, as so appearing, is hereby further amended by striking out, in the definition of "Covered model derivative", clause (iii) and inserting in place thereof the following clause:-

(iii) has been fine-tuned using a quantity of computing power not exceeding a threshold determined by the secretary pursuant to section 7, the cost of which, as reasonably assessed by the developer, exceeds $10,000,000 if calculated using the average market price of cloud compute at the start of fine-tuning;.

SECTION 5. Subsection (e) of section 2 of said chapter 93M, as so appearing, shall take effect on January 1, 2026.

SECTION 6. Sections 3 and 4 shall take effect on January 1, 2027.