STATE OF NEW YORK

10731

IN SENATE

October 28, 2026

Introduced by Sens. GOUNARDES, FAHY -- read twice and ordered printed, and when printed to be committed to the Committee on Rules

AN ACT to amend the general business law, in relation to the responsible AI safety and education (RAISE) act

The People of the State of New York, represented in Senate and Assembly, do enact as follows:

Section 1. Section 1420 of the general business law, as added by chapter 96 of the laws of 2026, is amended to read as follows:

§ 1420. Definitions. As used in this article, the following terms shall have the following meanings:

1. "Affiliate" means a person controlling, controlled by, or under common control with a specified person, directly or indirectly, through one or more intermediaries.

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

3. (a) "Catastrophic risk" means a foreseeable and material risk that a frontier developer's development, storage, use, or deployment of a frontier model will materially contribute to the death of, or serious injury to, more than fifty people or more than one billion dollars in damage to, or loss of, property arising from a single incident **or multiple incidents** involving a frontier model doing any of the following:

(i) providing expert-level assistance in the creation or release of a chemical, biological, radiological, or nuclear weapon;

(ii) engaging in conduct with [~~no meaningful~~] **limited** human oversight, intervention, or supervision that is either a cyberattack or, if the conduct had been committed by a human, would constitute the crime of murder, assault, extortion, or theft, including theft by false pretense; [~~or~~]

(iii) **attempting to overthrow or seize control of the federal government or state or local governments in the state of New York; or**

(iv) evading the control of its frontier developer or user.

EXPLANATION--Matter in **italics** (underscored) is new; matter in brackets [-] is old law to be omitted.

LBD16516-01-6

(b) "Catastrophic risk" does not include a foreseeable and material risk from any of the following:

(i) information that a frontier model outputs if the information is otherwise publicly accessible in a substantially similar form from a source other than a foundation model;

(ii) lawful activity of the federal government; or

(iii) harm caused by a frontier model in combination with other software if the frontier model did not materially contribute to the harm.

4. **"Covered emergency" means a situation that poses an imminent catastrophic risk or security risk. "Covered emergency" shall be construed to include a situation in which immediate action is necessary to prevent a catastrophic risk or security risk, even if the harm that would be prevented is not proximate in time to the situation.**

5. "Critical safety incident" means any of the following:

(a) unauthorized access to, modification of, or exfiltration of[~~, the~~] **unreleased** model weights of a frontier model [~~that results in death or bodily injury~~];

(b) harm resulting from the materialization of a catastrophic risk;

(c) [~~loss of control of a frontier model causing death or bodily injury; or~~

~~(d) a frontier model that uses deceptive techniques against the frontier developer to subvert the controls or monitoring of its frontier developer outside of the context of an evaluation designed to elicit this behavior and in a manner that demonstrates materially increased catastrophic risk~~] **conduct by a foundation model of a frontier developer that circumvents, disables, evades, or attempts to circumvent, disable, or evade, a material, technical or organizational control of its frontier developer that limits the model's access to systems, resources, or data, provides for its suspension or modification, or monitors its conduct or capabilities, outside of the context of an evaluation designed to elicit such conduct where such conduct remains confined to the evaluation environment;**

(d) conduct by a frontier model that, without the intent of any user, materially impairs the confidentiality, integrity, or availability of a computer system or computer network which is not owned or operated by the frontier developer or user and where such conduct occurs without the consent of the owner or operator of such system or network;

(e) an occurrence in which a safeguard against an incident described in this subdivision failed, was circumvented, or was materially degraded, such that, absent timely intervention, mitigation, or fortuity, there was a substantial likelihood that such a critical safety incident would have resulted; or

(f) any other occurrence or condition designated by regulation of the office which is related to catastrophic risk, security risk, or the automation of AI research and development.

6. "Department" means the department of financial services.

[~~5.~~] **7.** (a) "Deploy" means to make a frontier model available to a third party for use, modification, copying, or combination with other software.

(b) "Deploy" does not include making a frontier model available to a third party for the primary purpose of developing or evaluating the frontier model.

[~~6.~~] **8.** "Foundation model" means an artificial intelligence model that is all of the following:

(a) trained on a broad data set;

(b) designed for generality of output; and

(c) adaptable to a wide range of distinctive tasks.

[~~7.~~] 9. "Frontier AI framework" means documented technical and organizational protocols to:

(a) manage, assess, and mitigate catastrophic risks; and

(b) assess the ability of frontier models to automate AI research and development and to manage any catastrophic risks and security risks associated with such automation.

[~~8.~~] 10. "Frontier developer" means a person who has trained, or initiated the training of, a frontier model, with respect to which the person has used, or intends to use, at least as much computing power to train the frontier model as would meet the technical specifications found in subdivision [~~nine~~] twelve of this section.

11. "Frontier examiner" means a third party that performs services under section fourteen hundred twenty-four of this article and that is eligible to be a frontier examiner under such section.

[~~9.~~] 12. (a) "Frontier model" means a foundation model that was trained using a quantity of computing power greater than 10°[~~26~~] 25 integer or floating-point operations, provided, however, that foundation models shall be considered to be frontier models if the office has determined that the foundation model:

(i) is among the most capable foundation models;

(ii) possesses capabilities which pose catastrophic risk; and

(iii) is reasonably likely to possess the ability to automate AI research and development.

(b) The quantity of computing power described in paragraph (a) of this subdivision shall include computing for the original training run and for any subsequent fine-tuning, reinforcement learning, or other material modifications the developer applies to a preceding foundation model.

13. "Internal use" means any use of a frontier model by its frontier developer or its affiliates, including, but not limited to, any of the following:

(a) the training of the frontier model;

(b) modifying or fine-tuning the frontier model;

(c) using the frontier model for internal business purposes;

(d) evaluating the frontier model; and

(e) using the frontier model for AI research and development.

[~~10.~~] 14. "Large frontier developer" means a frontier developer that together with its affiliates collectively [~~had annual gross revenues in excess of five hundred million dollars in the preceding calendar year~~] incurred more than one billion dollars in expenditures for research, development, training, fine-tuning, modification, security, evaluation, or other activities which have contributed to the development, creation, or modification of one or more frontier models in the preceding thirty-six months.

[~~11.~~] 15. "Model weight" means a numerical parameter in a frontier model that is adjusted through training and that helps determine how inputs are transformed into outputs.

[~~12. "Department" means the department of financial services.~~

~~13. "Property" means tangible or intangible property.~~

~~14. "Person" means 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.~~

~~15.~~] 16. "Office" means an office within the department of financial services, which shall report to the superintendent of financial services and is tasked with implementation of this article.

17. "Open source release" means making the model weights of a foundation model publicly available under a license that grants to all persons, without payment of a fee or royalty and without restriction or discrimination based on the identity of any person or group of persons or on the field of endeavor, the rights to access, use, modify, and redistribute the foundation model or modified versions of the foundation model. A license does not fail to satisfy this subdivision solely because it requires attribution to the frontier developer or requires that modified or redistributed versions be made available under identical or comparable terms.

18. "Person" means 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.

19. "Property" means tangible or intangible property.

20. "Relevant information" means all information reasonably necessary for a frontier examiner to comply with section fourteen hundred twenty-four of this article, including, but not limited to:

(a) all unredacted versions of materials published pursuant to this article;

(b) access to frontier models sufficient to enable assessments under section fourteen hundred twenty-four of this article;

(c) interviews with employees of a large frontier developer;

(d) internal documents of a large frontier developer; and

(e) information to which an employee of a large frontier developer whose responsibilities include assessing the capabilities of frontier models or managing catastrophic risk or security risk has access in the ordinary course of employment is relevant information unless the large frontier developer establishes that it is not reasonably necessary for the frontier examiner to perform its functions under section fourteen hundred twenty-four of this article.

21. "Specification" means a high-level, general description of the principles that the model is intended to follow, how the model is intended to prioritize different principles and user prompts, and the topics for which the model is intended to refuse, redirect, or otherwise subvert user prompts.

22. "Security risk" means a foreseeable and material risk of unauthorized access to, modification of, or transfer of the unreleased model weights of a frontier model by any person, entity, or model.

23. "Superintendent" means the superintendent of financial services.

[~~16. "Office" means an office within the department of financial services, which shall report to the superintendent of financial services and is tasked with implementation of this article.~~]

§ 2. Section 1421 of the general business law, as added by chapter 96 of the laws of 2026, is amended to read as follows:

§ 1421. Transparency requirements. 1. A large frontier developer shall write, implement, comply with, and clearly and conspicuously publish on its internet website a frontier AI framework that applies to the large frontier developer's frontier models and describes in detail how the large frontier developer handles all of the following:

(a) incorporating national standards, international standards, and industry consensus best practices into its frontier AI framework;

(b) defining and assessing thresholds used by the large frontier developer to identify and assess whether a frontier model has capabilities that could pose a catastrophic risk or a security risk or whether a

frontier model has the capability to automate AI research and development, which may include multiple-tiered thresholds;

(c) applying mitigations to address the potential for catastrophic or security risks based on the results of assessments undertaken pursuant to paragraph (b) of this subdivision;

(d) reviewing assessments and adequacy of mitigations as part of the decision to deploy a frontier model or use it [~~extensively~~] internally;

(e) [~~using third parties to assess the potential for catastrophic risks and the effectiveness of mitigations of catastrophic risks;~~

~~(f)~~] revisiting and updating the frontier AI framework, including any criteria that trigger updates and how the large frontier developer determines when its frontier models are substantially modified enough to require disclosures pursuant to subdivision [~~three~~] four of this section;

[~~(g)~~] (f) cybersecurity and information security practices to [~~secure unreleased model weights from unauthorized modification or transfer by internal or external parties~~] mitigate security risks based on the results of assessments undertaken pursuant to paragraph (b) of this subdivision;

[~~(h)~~] (g) identifying and responding to critical safety incidents;

[~~(i)~~] (h) instituting internal governance practices to ensure implementation of [~~these processes~~] the frontier AI frameworks; [~~and~~

~~(j)~~] (i) assessing and managing catastrophic risk resulting from the internal use of its frontier models, including risks resulting from a frontier model circumventing oversight mechanisms and risks related to the frontier model being used to automate AI research and development;

(j) assessing frontier models for signs of exploiting, undermining, or evading the large frontier developer's training, assessment, or mitigation processes or for deceptive or manipulative behavior; and

(k) whether and how the large frontier developer uses frontier models to monitor and evaluate other frontier models.

2. The office may adopt regulations to add additional elements that a large frontier developer is required to include in its frontier AI framework to account for changes in the scientific understanding of catastrophic risk, changes in AI capabilities or mitigation approaches, or to add additional elements that would effectuate the purposes of this chapter.

3. (a) A large frontier developer shall review and, as appropriate, update its frontier AI framework at least once per year.

(b) If a large frontier developer makes a material modification to its frontier AI framework, the large frontier developer shall clearly and conspicuously publish the modified frontier AI framework and a justification for that modification within thirty days.

[~~3.~~] 4. (a) Before, or concurrently with, deploying a new frontier model or a substantially modified version of an existing frontier model, a frontier developer shall clearly and conspicuously publish on its internet website a transparency report containing all of the following:

(i) the internet website of the frontier developer;

(ii) a mechanism that enables a natural person to communicate with the frontier developer;

(iii) the release date of the frontier model;

(iv) the languages supported by the frontier model;

(v) the modalities of output supported by the frontier model;

(vi) the intended uses of the frontier model; [~~and~~]

(vii) the frontier model specification; and

(viii) any generally applicable restrictions or conditions on uses of the frontier model.

(b) Before, or concurrently with, deploying a new frontier model or a substantially modified version of an existing frontier model, a large frontier developer shall include in the transparency report required by paragraph (a) of this subdivision, summaries of all of the following:

(i) assessments of catastrophic risks or security risks from the frontier model conducted pursuant to the large frontier developer's frontier AI framework;

(ii) the results of [~~the~~] such assessments [~~under subparagraph (i) of this paragraph~~];

(iii) the extent to which third-party evaluators were involved in such assessments; [~~and~~]

(iv) the extent to which such assessments are reproducible by external entities; and

(v) other steps taken to fulfill the requirements of the frontier AI framework with respect to the frontier model.

(c) A frontier developer that publishes the information described in paragraph (a) or (b) of this subdivision as part of a larger document, including a system card or model card, shall be deemed in compliance with the applicable paragraph.

[~~4.~~] 5. (a) (i) A frontier developer shall not make a materially false or misleading statement about catastrophic risk or security risk from its frontier models or its management of catastrophic risk or security risk.

(ii) A large frontier developer shall not make a materially false or misleading statement about its implementation of, or compliance with, its frontier AI framework.

(iii) A frontier developer shall not make a materially false or misleading statement in any publication, submission, or filing to the office under this article.

(iv) A large frontier developer shall designate an accountable senior officer responsible for compliance with this chapter. The accountable senior officer shall certify, after reasonable inquiry and under penalty of perjury, that, to the best of the officer's knowledge, each publication and material submission to the office is accurate, complete, and not materially misleading.

(b) [~~This subdivision~~] Subparagraphs (i), (ii) and (iii) of paragraph (a) of this subdivision shall not apply to a statement that was made in good faith and was reasonable under the circumstances.

[~~5.~~] 6. (a) A large frontier developer shall implement technical and administrative measures that are sufficient to allow it to do both of the following:

(i) prevent an identified imminent critical safety incident; and

(ii) contain a critical safety incident as defined in subdivision five of section fourteen hundred twenty of this article.

(b) This subdivision shall not be construed to require technical and administrative measures sufficient to prevent or contain a critical safety incident that arises from a copy of a frontier model that has been distributed in an open source release and that is not in the possession of its frontier developer.

7. (a) When a frontier developer publishes documents to comply with this section, such frontier developer may make redactions to such documents that are necessary to protect such frontier developer's trade secrets, such frontier developer's cybersecurity, public safety, or the

national security of the United States or to comply with any federal or state law.

(b) If a frontier developer redacts information in a document pursuant to this subdivision, such frontier developer shall:

(i) describe the character and justification of such redaction in any published version of such document to the extent permitted by the concerns that justify redaction [~~and shall~~];

(ii) retain the unredacted information for five years; and

(iii) grant the office access to an unredacted copy of the publication upon request.

§ 3. Section 1422 of the general business law, as added by chapter 96 of the laws of 2026, is renumbered to be section 1423 and a new section 1422 is added to read as follows:

§ 1422. Minimum standards. 1. No later than January first, two thousand twenty-nine, the office shall adopt regulations, which shall be binding on large frontier developers, establishing minimum standards for the frontier AI frameworks required in section fourteen hundred twenty-one of this article as well as minimum standards for required mitigations of catastrophic risk and security risk which shall address:

(a) each element required to be included in subdivision one of section fourteen hundred twenty-one of this article;

(b) each type of catastrophic risk enumerated in the definition of catastrophic risk in subdivision three of section fourteen hundred twenty of this article;

(c) security risks; and

(d) the automation of AI research and development.

2. In implementing this section, the office shall aim to promulgate regulations that, if successfully implemented, would come as close as practicable to preventing catastrophic risks and security risks that are unreasonable in light of the societal benefits of the activities that create those risks.

3. In implementing this section, the office shall consider whether, given the current state of knowledge regarding AI safety and security, the subject matter of the proposed regulation is sufficiently understood to support effective regulation, or whether the purposes of this article would be better served through voluntary guidance, monitoring, or further study.

4. In implementing this section, the office shall consider relevant national standards, international standards, and industry consensus best practices, including existing frontier AI frameworks published by frontier developers.

5. The office shall consult with frontier developers, frontier examiners, and other technical experts when creating the minimum standards prescribed herein.

§ 4. Section 1423 of the general business law, as added by chapter 96 of the laws of 2026, and as renumbered by section three of this act, is amended to read as follows:

§ 1423. Reporting. 1. The office shall establish a mechanism to be used by a frontier developer or a member of the public to report a critical safety incident that includes all of the following:

(a) the date of the critical safety incident;

(b) the reasons the incident qualifies as a critical safety incident;

(c) a short and plain statement describing the critical safety incident; [~~and~~]

(d) whether the incident was associated with internal use of a frontier model; and

(e) any other information the office may deem necessary to better understand such incident or the state of catastrophic risk, security risk, and the automation of AI research and development.

2. (a) A large frontier developer shall transmit to the office a summary of any assessment of:

(i) catastrophic risk resulting from internal use of its frontier models [~~every three months or pursuant to another reasonable schedule requested by the large frontier developer, communicated in writing to the office with written updates, as appropriate, and agreed upon by the office.~~];

(ii) the ability of its frontier models to automate AI research and development and their use in AI research and development, including a specific percentage of how much of such research and development is currently automated;

(iii) the percentage of computing power used for models that are internally used and/or deployed; and

(iv) a general description of how such models are being internally used and/or deployed.

(b) A large frontier developer shall transmit the assessments required under paragraph (a) of this subdivision at all of the following times:

(i) every three months or pursuant to another reasonable schedule requested by the large frontier developer, communicated in writing to the office, and agreed upon by the office;

(ii) within seventy-two hours of any time that there has been a substantial increase in the large frontier developer's assessment of catastrophic risks or security risks from internal use of its frontier models, or capabilities of its frontier models to automate AI research and development, in comparison to the previously reported assessment; and

(iii) at any other time that the office may by regulation require to ensure the office can adequately monitor the state of catastrophic risks, security risks, the ability of frontier models to automate AI research and development, and frontier models' use in AI research and development.

The office shall establish a mechanism to be used by a large frontier developer to confidentially submit summaries of any assessments [~~of the potential for catastrophic risk resulting from internal use of its frontier models~~] required under paragraph (a) of this subdivision, provided, however, that such summaries shall be shared with the attorney general. The attorney general shall be empowered to pursue the actions provided in subdivision two of section fourteen hundred thirty of this article as set forth in such subdivision.

[~~(b)~~] (c) The office shall take all reasonable precautions to limit access to any reports related to internal use of frontier models [~~to only personnel authorized to know the information~~] and to protect the reports from unauthorized access.

3. (a) [~~Subject to paragraph (b) of this subdivision, a~~] A frontier developer shall report any critical safety incident pertaining to one or more of its frontier models to the office within seventy-two hours from a determination that a critical safety incident has occurred or within seventy-two hours of the frontier developer learning facts sufficient to establish a reasonable belief that a critical safety incident has occurred.

(b) If a frontier developer discovers that a critical safety incident poses an imminent risk of death or serious physical injury, the frontier developer shall disclose that incident within twenty-four hours to an

authority, including any law enforcement agency or public safety agency with jurisdiction, that is appropriate based on the nature of that incident and as required by law.

(c) A frontier developer that discovers information about a critical safety incident after filing the initial report required by this subdivision may file an amended report.

4. The office shall review**:**

(a) critical safety incident reports submitted by frontier developers [~~and~~]**;**

(b) assessments transmitted by frontier developers under subdivision two of this section; and

(c) reports, notifications, and other information transmitted to the office by a frontier examiner pursuant to this article.

The office may review reports submitted by members of the public.

5. (a) The office may transmit reports of critical safety incidents or summaries of any assessments [~~of catastrophic risk from internal use of frontier models~~] **required under subdivision two of this section** to other governmental entities at their discretion, considering for example and without limitation the following: the severity of any such incident, potential ongoing risks, legal or regulatory obligations, the need for coordinating with other governmental agencies or other entities and the availability of information. [~~The office shall consider transmitting such reports or summaries to the office of the attorney general, as appropriate.~~] Any report transmitted from the office to another governmental entity shall be exempt from disclosure under article six of the public officers law.

(b) The office may consider, at its discretion, any risks related to trade secrets, public safety, cybersecurity of a frontier developer, or national security when transmitting reports.

6. A report of a critical safety incident submitted to the office pursuant to this section and a report of assessments of catastrophic risk from internal use pursuant to section fourteen hundred twenty-one of this article, are exempt from disclosure under article six of the public officers law.

7. (a) Beginning January first, two thousand twenty-eight, and annually thereafter, the office shall produce a report, that includes the following:

(i) anonymized and aggregated information about critical safety incidents that have been reviewed by the office since the preceding report;

(ii) any information that the office deems relevant to frontier model safety;

(iii) recommended updates to this article, if any; and

(iv) any developments relevant to the purposes of this article.

(b) The office shall not include information in a report pursuant to this subdivision that would compromise the trade secrets or cybersecurity of a frontier developer, public safety, or the national security of the United States or that would be prohibited by any federal or state law.

(c) The office shall transmit a report pursuant to this subdivision to the governor, the temporary president and minority leader of the senate, the speaker and minority leader of the assembly, the chair and ranking member of the senate committee on internet and technology, and the chair and ranking member of the assembly committee on science and technology.

8. The office may adopt regulations designating one or more federal laws, regulations, or guidance documents that meet all of the following conditions for the purposes of subdivision nine of this section:

(a) (i) the law, regulation, or guidance document imposes or states standards or requirements for critical safety incident reporting that are substantially equivalent to, or stricter than, those required by subdivision three of this section; and

(ii) the law, regulation, or guidance document described in subparagraph (i) of this paragraph does not need to require critical safety incident reporting to the state of New York; and

(b) the law, regulation, or guidance document is intended to assess, detect, or mitigate the catastrophic risk.

9. (a) A frontier developer that intends to comply with subdivision three of this section by complying with the requirements of, or meeting the standards stated by, a federal law, regulation, or guidance document designated pursuant to subdivision eight of this section shall declare its intent to do so to the office.

(b) After a frontier developer has declared its intent pursuant to paragraph (a) of this subdivision, the following shall apply:

(i) the frontier developer shall be deemed in compliance with subdivision three of this section to the extent that the frontier developer meets the standards of, or complies with the requirements imposed or stated by, the designated federal law, regulation, or guidance document until the frontier developer declares the revocation of that intent to the office or the office revokes a relevant regulation pursuant to subdivision ten of this section;

(ii) the failure by a frontier developer to meet the standards of, or comply with the requirements stated by, the federal law, regulation, or guidance document designated pursuant to subdivision eight of this section shall constitute a violation of this article; and

(iii) frontier developers who comply with subdivision three of this section by meeting such federal standards shall send copies of any critical safety incident reports required by such federal standards to the office concurrently with sending them to federal authorities.

10. The office shall revoke a regulation adopted under subdivision eight of this section if the requirements of subdivision eight of this section are no longer met.

§ 5. Sections 1423, 1424, 1425, 1426, 1427, 1428 and 1429 of the general business law, as added by chapter 96 of the laws of 2026 are renumbered sections 1427, 1428, 1429, 1430, 1431, 1432 and 1433, respectively, and three new sections 1424, 1425 and 1426 are added to read as follows:

§ 1424. Frontier examiners. 1. (a) No later than sixty days after the effective date of this article or sixty days after the date on which a person first qualifies as a large frontier developer, whichever is later, and thereafter on an annual basis, a large frontier developer shall retain a frontier examiner to produce a report which assesses the following:

(i) whether a large frontier developer has complied with its frontier AI framework as well as any instances of noncompliance with such framework;

(ii) whether such AI framework may require additional detail in order to better assess such compliance; and

(iii) whether any statements made by the large frontier developer to the public, office, attorney general, or other governmental entities regarding catastrophic risks or security risks, such developer's management of such risks, and/or such developer's compliance with its frontier AI framework are inconsistent with the findings of the frontier examiner's report.

(b) A frontier examiner shall transmit an unredacted version of the final report to the office and the large frontier developer within seven days of completing such report; provided, however, that such unredacted report shall be exempt from disclosure under article six of the public officers law.

(c) No less than twenty-one days after transmitting the final report under paragraph (b) of this subdivision, the frontier examiner shall conspicuously publish such report with appropriate redactions as provided in paragraph (e) of this subdivision.

(d) Within seven days of a frontier examiner publishing a report pursuant to paragraph (c) of this subdivision, the large frontier developer shall conspicuously publish a link to such report on such developer's public website.

(e) Any public version of a frontier examiner's report under this subdivision may contain redactions to the extent necessary to protect the frontier examiner or large frontier developer's trade secrets, the frontier examiner's or frontier developer's cybersecurity, public safety, or the national security of the United States, or to comply with any federal or state law.

2. (a) On or before January first, two thousand twenty-eight, a large frontier developer shall retain a frontier examiner to perform ongoing, continuous assessment of the large frontier developer for the purposes of producing the outputs required by this subdivision.

(b) On or before April first, two thousand twenty-eight, and every ninety days thereafter, a frontier examiner retained by a larger frontier developer under this subdivision shall produce a report that includes:

(i) an assessment of each type of catastrophic risk posed by the large frontier developer's frontier models, taking into account model capabilities, applicable threat models, and the large frontier developer's mitigations;

(ii) an assessment of security risks related to the large frontier developer's frontier model;

(iii) an assessment of the ability of the large frontier developer's frontier models to automate artificial intelligence research and development and any increased potential for catastrophic risks or challenges to risk monitoring, assessment or mitigation resulting from such ability;

(iv) an assessment of whether the large frontier developer's activities would pose a substantial level of catastrophic risk over the next ninety day period, if additional mitigations were not put in place; and

(v) recommended mitigations, if any, that would reduce the level of catastrophic risk.

(c) A frontier examiner shall further include in an appendix to the report described in this subdivision:

(i) a list of the personnel involved in the report and the processes involved in creating such report;

(ii) the frontier examiner's procedures for managing conflicts of interest and any specific conflicts of interest of any personnel involved in the report;

(iii) the methodology of the process used to produce the report and the nature of the information reviewed by the frontier examiner to produce the report; and

(iv) the signature of the lead examiner certifying the results of the report.

(d) A frontier examiner shall transmit an unredacted version of the final report and its appendix to the office and the large frontier developer within seven days of completing such report, provided, however, that such unredacted report shall be exempt from disclosure under article six of the public officers law.

(e) No less than twenty-one days after transmitting the final report under paragraph (d) of this subdivision, the frontier examiner shall conspicuously publish such report with appropriate redactions as provided in paragraph (g) of this subdivision.

(f) Within seven days of a frontier examiner publishing a report pursuant to paragraph (e) of this subdivision, the large frontier developer shall conspicuously publish a link to such report on such developer's public website.

(g) Any public version of a frontier examiner's report under this subdivision may contain redactions to the extent necessary to protect the frontier examiner or large frontier developer's trade secrets, the frontier examiner's or frontier developer's cybersecurity, public safety, or the national security of the United States, or to comply with any federal or state law.

(h) If a redaction is made pursuant to this subdivision, the person making or requesting the redaction shall provide a written statement with the character and justification of the redaction that is included in any published version of the document to the extent permitted by the concerns that justify redaction, and shall retain the unredacted information for five years.

(i) If a large frontier developer requests a redaction and a frontier examiner believes that such redaction does not comply with paragraph (g) of this subdivision, the frontier examiner shall notify the office in writing within fourteen days of publishing the examiner's report. Upon receipt of such notification or upon the office's own determination that a redaction is not in compliance with this subdivision the office shall:

(i) notify the large frontier developer and provide the large frontier developer at least ten days to submit a written justification for such redaction; and

(ii) within fifteen days of receiving such justification from the developer, determine in writing whether such redaction complies with paragraph (g) of this subdivision and provide a copy of such determination to the frontier examiner and large frontier developer.

(j) If the office determines pursuant to paragraph (i) of this subdivision that redaction is not in compliance with paragraph (g) of this subdivision, the frontier examiner shall publish its report and appendix without such redaction within ten days of such office's written determination.

(k) A large frontier developer shall grant its frontier examiner access to all relevant information as defined in subdivision twenty of section fourteen hundred twenty of this article.

(l) To protect a large frontier developer's trade secrets, confidential business information, cybersecurity, national security of the United States, or public safety, a large frontier developer may impose reasonable security protocols on the frontier examiner, including, but not limited to restrictions on note taking, copying, retaining, or removing materials, requirements for on-premise review, and confidentiality requirements.

(m) Notwithstanding paragraph (l) of this subdivision, a large frontier developer shall not impose security protocols on the frontier examiner that would prevent the frontier examiner from reasonably carrying

out its responsibilities under this section, and a frontier examiner which believes such security protocols have been imposed shall notify the large frontier developer and the office in writing.

(n) Upon receipt of a notification pursuant to paragraph (m) of this subdivision, the office shall:

(i) provide the developer at least ten days to either submit a written justification for withholding access or information from the frontier examiner, or to provide such access or information;

(ii) determine within ten days of receiving the information in subparagraph (i) of this paragraph whether information or access is required to be provided by the developer under this section; and

(iii) if the office determines that such information or access is so required, provide written notice of such determination to the frontier examiner and large frontier developer.

(o) The large frontier developer shall provide such information or access required under subparagraph (iii) of paragraph (n) of this subdivision to the frontier examiner within ten days of receiving notice of such requirement from the office.

(p) A large frontier developer shall retain for five years each document created, published, or submitted pursuant to this article alongside records sufficient to permit a frontier examiner to verify: the developer's assessments of frontier models; the developer's decisions to deploy or use frontier models, and the developer's handling of critical safety incidents. Upon written request by the frontier examiner or the office, a large frontier developer shall preserve all relevant information identified in such request for five years.

3. (a) Any entity that seeks to perform the duties of a frontier examiner under this article shall:

(i) employ or retain personnel for its duties as a frontier examiner who have technical expertise in catastrophic risks, security risks, and the automation of AI research and development;

(ii) have and implement policies to manage conflicts of interest that could impair the independence or integrity of the examiner's work; and

(iii) not have financial, operational, or managerial dependence on any large frontier developer or affiliate of a large frontier developer; provided, however, that reasonable compensation for services performed as a frontier examiner pursuant to this article shall not constitute financial dependence for purposes of this subparagraph if the compensation is not contingent on the content of any report, finding, or recommendation of the frontier examiner.

(b) Beginning January first, two thousand twenty-nine, only entities licensed as a frontier examiner under subdivision four of this section shall perform the duties of a frontier examiner under this article.

4. (a) No later than January first two thousand twenty-eight, the office shall adopt regulations establishing a procedure for the licensing of frontier examiners who meet the requirements of paragraph (a) of subdivision three of this section and any other requirements the office determines would further the purposes of this article.

(b) Such regulations shall detail the procedure by which an entity may apply to be a frontier developer and the procedure by which the office shall approve or deny such application within ninety days.

(c) The regulations shall detail conditions under which and the procedure by which the office shall revoke the license of a frontier developer.

(d) The regulations shall ensure the integrity of examinations conducted by licensed examiners by:

(i) including conflict of interest rules for a frontier examiner, its employees, and its contractors; and

(ii) including data security and handling rules for frontier examiners.

(e) The office may in its discretion require a large frontier developer to retain an additional specified frontier examiner:

(i) In cases where a critical safety incident has been reported to the office and a frontier examiner is needed to investigate such critical safety incident; and/or

(ii) In cases where the office finds that a large frontier developer's original frontier examiner has not performed its duties in a manner that permits the office to have satisfactory visibility into the level of catastrophic risk and security risk posed by the large frontier developer's activities.

(f) When requiring a large frontier developer to retain an additional frontier developer under paragraph (e) of this subdivision, the office shall describe in detail the identity of the designated additional frontier examiner, the reasons for such designation, the form and timing of reports that the additional frontier examiner is expected to produce in the event of a critical safety incident investigation, and the duration of such additional examination, which shall expire after one year absent redesignation by the office.

5. A large frontier developer shall not retain an entity as a frontier examiner if the entity is not eligible to be a frontier examiner under this section. Beginning January first, two thousand twenty-nine, a large frontier developer may only retain entities licensed as frontier examiners under subdivision four of this section.

6. Nothing in this section shall authorize the imposition of a civil penalty solely for a large frontier developer's failure to adopt a recommendation of a frontier examiner where no court order requiring the recommended action has been issued in an action brought pursuant to section fourteen hundred thirty-one of this article.

7. (a) If, at any time, a frontier examiner believes that there is a covered emergency related to a large frontier developer's activities, the frontier examiner shall submit a report to the large frontier developer, the office, and the attorney general as soon as practicable, but no later than seventy-two hours after forming such belief.

(b) The frontier examiner shall also submit a report containing the matters required by this subdivision to the large frontier developer, the office, and the attorney general within seventy-two hours after the attorney general or the office requests such a report in connection with an action brought or under consideration pursuant to section fourteen hundred thirty-one of this article, whether or not the frontier examiner has independently formed the belief described in paragraph (a) of this subdivision.

(c) A report under this subdivision shall state the frontier examiner's assessment of whether there is a covered emergency, the factual and technical basis for that assessment, any recommendations for actions the large frontier developer could take to prevent or materially mitigate the covered emergency, and any material information the frontier examiner was unable to obtain.

§ 1425. Computing power limitations. 1. The office shall be empowered to adopt regulations prescribing the circumstances under which, and the procedure by which, it may order that large frontier developers limit computing power as described in this section. Such order shall apply to all large frontier developers simultaneously, provided however that the

office shall retain the discretion to order different limitations of computing power per large frontier developer as the office deems necessary and in order to achieve the objective of subdivision two of this section.

2. Such order may only be issued in order to reduce catastrophic risk, prevent a critical safety incident, or avert a covered emergency. Such order may be informed by reports or assessments submitted or published under this article or by information otherwise available to the office.

3. Such order may limit the amount of computing power used to train new models of a large frontier developer and/or the amount of computing power dedicated to automated research and development.

4. The office shall retain a frontier examiner as described in section fourteen hundred twenty-four of this article to verify compliance with such order.

§ 1426. Whistleblower protections. 1. A large frontier developer who is an employer shall be bound by all of the provisions of section seven hundred forty of the labor law.

2. A large frontier developer shall not make, adopt, enforce, or enter into any rule, regulation, policy, or contract that prevents an employee from engaging in any of the protected activities provided in paragraph (a), (b) or (c) of subdivision two of section seven hundred forty of the labor law.

3. An employee may use the website created under subdivision three of section fourteen hundred thirty-one of this article to disclose any activity under this section.

4. (a) In addition to the notice requirements provided in subdivision eight of section seven hundred forty of the labor law, a large frontier developer shall ensure that any new employees receive equivalent notice of such employee's protections, rights, and obligations under such section within ten days of the initial date of such employee's employment.

(b) A large frontier developer shall ensure that employees who work remotely receive equivalent notice of their protections, rights, and obligations under section seven hundred forty of the labor law.

(c) A large frontier developer shall provide written notice to each covered employee of such employee's rights under section seven hundred forty of the labor law and under this section at least once each year. A large frontier developer shall ensure that the notice is received and acknowledged by all employees.

§ 6. Section 1431 of the general business law, as added by chapter 96 of the laws of 2026 and as renumbered by section five of this act, is amended to read as follows:

§ 1431. Violations. 1. The attorney general may bring a civil action to recover a civil penalty in an amount not to exceed one million dollars [~~for a first violation~~] per day that a first violation occurs and in an amount not to exceed three million dollars per [~~subsequent violation, determined based on the severity of the violation~~] day that a subsequent violation occurs where a large frontier developer fails to [~~publish or transmit a compliant document required to be published or transmitted under this article, makes a statement in violation of subdivision four of section fourteen hundred twenty-one of this article, fails to report an incident as required by section fourteen hundred twenty-two of this article, or fails to comply with its own frontier AI framework~~] comply with this chapter, a regulation adopted pursuant to this article, or a court order issued pursuant to this section.

2. (a) The attorney general may further bring an action in a court of competent jurisdiction for injunctive relief, including a temporary restraining order, against a large frontier developer where:

(i) The large frontier developer violates this article or fails to comply with a regulation adopted pursuant to this article;

(ii) A frontier examiner reports pursuant to subdivision seven of section fourteen hundred twenty-four of this article that there is a covered emergency related to the large frontier developer's activities; or

(iii) The attorney general otherwise has reason to believe, including on the basis of a report or assessment submitted or published under this article, that there is a covered emergency related to the large frontier developer's activities.

(b) In an action under this section, the court may order any relief reasonably necessary to prevent or materially mitigate the covered emergency or remedy the violation, including corrective actions, the cessation of any activity, and/or an order restricting the deployment, use, or availability of a model within the state. The court may modify or dissolve an order as circumstances change.

(c) The court shall preserve the confidentiality of information exempt from disclosure under this article and/or redacted under this article by reasonable means, including by sealing portions of the record, while making its decision and the basis for the decision public to the maximum extent consistent with that protection.

(d) The office shall promptly refer to the attorney general any information indicating a violation of this chapter or a covered emergency, shall provide to the attorney general on request any record submitted to the office under this article, and shall provide technical assistance in an action under this section at the attorney general's request.

3. The attorney general shall maintain a website to receive complaints, information, and/or referrals from persons regarding an operator's or covered platform's alleged compliance or noncompliance with the provisions of this article.

4. Nothing in this article shall be construed to establish, authorize or create a private right of action associated with violations of this article.

[~~3.~~] 5. Nothing in this article shall be construed to prevent a large frontier developer from asserting that another person, entity, or factor, may be responsible for any alleged harm, injury or damage resulting from a catastrophic risk or critical safety incident.

§ 7. Severability. If any clause, sentence, paragraph, subdivision, section or part of this act shall be adjudged by any court of competent jurisdiction to be invalid, such judgment shall not affect, impair, or invalidate the remainder thereof, but shall be confined in its operation to the clause, sentence, paragraph, subdivision, section, or part thereof directly involved in the controversy in which such judgment shall have been made.

§ 8. This act shall take effect on the on the same date and in the same manner as chapter 699 of the laws of 2025, takes effect.