24105785D

HOUSE BILL NO. 747

AMENDMENT IN THE NATURE OF A SUBSTITUTE

(Proposed by the House Committee on Communications, Technology and Innovation on February 5, 2024)

(Patron Prior to Substitute—Delegate Maldonado)

A BILL to amend the Code of Virginia by adding in Title 59.1 a chapter numbered 57, consisting of sections numbered 59.1-603 through 59.1-607, relating to High-risk Artificial Intelligence Developer Act established; civil penalty.

Be it enacted by the General Assembly of Virginia:

1. That the Code of Virginia is amended by adding in Title 59.1 a chapter numbered 57, consisting of sections numbered 59.1-603 through 59.1-607, as follows:

CHAPTER 57.

§ 59.1-603. Definitions.

As used in this chapter, unless the context requires a different meaning:

"Algorithmic discrimination" means any discrimination that is (i) prohibited under state or federal law and (ii) a reasonably foreseeable consequence of deploying or using a high-risk artificial intelligence system to make a consequential decision.

"Artificial intelligence" means technology that uses data to train statistical models for the purpose of enabling a computer system or service to autonomously perform any task, including visual perception, natural language processing, and speech recognition, that is normally associated with human intelligence or perception.

"Artificial intelligence system" means any software that incorporates artificial intelligence.

"Consequential decision" means any decision that has a material legal, or similarly significant, effect on a consumer's access to credit, criminal justice, education, employment, health care, housing, or insurance.

"Consumer" means a natural person who is a resident of the Commonwealth acting only in an individual or household context. "Consumer" does not include a natural person acting in a commercial or employment context.

"Deployer" means any person doing business in the Commonwealth that deploys or uses a high-risk artificial intelligence system to make a consequential decision.

"Developer" means any person doing business in the Commonwealth that develops or intentionally and substantially modifies a high-risk artificial intelligence system that is offered, sold, leased, given, or otherwise provided to consumers in the Commonwealth.

"Foundation model" means a machine learning model that (i) is trained on broad data at scale, (ii) is designed for generality of output, and (iii) can be adapted to a wide range of distinctive tasks.

"Generative artificial intelligence" means artificial intelligence based on a foundation model that is capable of and used to produce synthetic digital content, including audio, images, text, and videos.

"Generative artificial intelligence system" means any artificial intelligence system or service that incorporates generative artificial intelligence.

"High-risk artificial intelligence system" means any artificial intelligence system that is specifically intended to autonomously make, or be a controlling factor in making, a consequential decision. A system or service is not a "high-risk artificial intelligence system" if it is intended to (i) perform a narrow procedural task, (ii) improve the result of a previously completed human activity, (iii) detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment without proper human review, or (iv) perform a preparatory task to an assessment relevant to a consequential decision.

"Machine learning" means the development of algorithms to build data-derived statistical models that are capable of drawing inferences from previously unseen data without explicit human instruction.

"Significant update" means any new version, new release, or other update to a high-risk artificial intelligence system that results in significant changes to such high-risk artificial intelligence system's use case or key functionality.

"Synthetic digital content" means machine generated digital content, including any audio, image, text, or video that is produced by a generative artificial intelligence system.

§ 59.1-604. Operating standards for developers of high-risk artificial intelligence systems.

A. No developer of a high-risk artificial intelligence system shall offer, sell, lease, give, or otherwise provide to a deployer a high-risk artificial intelligence system unless the developer makes available to the deployer (i) a statement disclosing the intended uses of such high-risk artificial intelligence system and (ii) documentation disclosing (a) the known limitations of such high-risk artificial intelligence system, including any and all reasonably foreseeable risks of algorithmic discrimination arising from the

intended uses of such high-risk artificial intelligence system; (b) the purpose of such high-risk artificial intelligence system and the intended benefits and uses of such high-risk artificial intelligence system; (c) a summary describing how such high-risk artificial intelligence system was evaluated for performance and relevant information related to explainability before such high-risk artificial intelligence system was licensed or sold; (d) the measures the developer has taken to mitigate reasonable foreseeable risks of algorithmic discrimination that the developer knows arises from deployment or use of such high-risk artificial intelligence system; and (e) how an individual can use such high-risk artificial intelligence system to make, or monitor such high-risk artificial intelligence system when such high-risk artificial intelligence system is deployed or used to make, a consequential decision.

B. Each developer that offers, sells, leases, gives, or otherwise makes available to a deployer a high-risk artificial intelligence system shall make available to the deployer information and documentation in the developer's possession, custody, or control that is reasonably required to complete an impact assessment.

C. Nothing in this section shall be construed to require a developer to disclose any trade secret or confidential or proprietary information.

D. High-risk artificial intelligence systems that are in conformity with artificial intelligence standards, or parts thereof, shall be presumed to be in conformity with related requirements set out in this section and in associated regulations.

§ 59.1-605. Operating standards for deployers of high-risk artificial intelligence systems.

A. Each deployer of a high-risk artificial intelligence system shall avoid any risk of algorithmic discrimination that is a reasonably foreseeable consequence of deploying or using a high-risk artificial intelligence system to make a consequential decision.

B. No deployer shall deploy or use a high-risk artificial intelligence system to make a consequential decision unless the deployer has designed and implemented a risk management policy and program for such high-risk artificial intelligence system. The risk management policy shall specify the principles, processes, and personnel that the deployer shall use in maintaining the risk management program to identify, mitigate, and document any risk of algorithmic discrimination that is a reasonably foreseeable consequence of deploying or using such high-risk artificial intelligence system to make a consequential decision. Each risk management policy and program designed, implemented, and maintained pursuant to this subsection shall be (i) at least as stringent as the latest version of the Artificial Intelligence Risk Management Framework published by the National Institute of Standards and Technology or another nationally or internationally recognized risk management framework for artificial intelligence systems and (ii) reasonable considering (a) the size and complexity of the deployer; (b) the nature and scope of the high-risk artificial intelligence systems deployed and used by the deployer, including the intended uses of such high-risk artificial intelligence systems; (c) the sensitivity and volume of data processed in connection with the high-risk artificial intelligence systems deployed and used by the deployer; and (d) the cost to the deployer to implement and maintain such risk management program.

C. Except as provided in this subsection, no deployer shall deploy or use a high-risk artificial intelligence system to make a consequential decision unless the deployer has completed an impact assessment for such high-risk artificial intelligence system. The deployer shall complete an impact assessment for a high-risk artificial intelligence system (i) before the deployer initially deploys such high-risk artificial intelligence system and (ii) not later than 90 days after each significant update to such high-risk artificial intelligence system is made available.

Each impact assessment completed pursuant to this subsection shall include, at a minimum:

1. A statement by the deployer disclosing (i) the purpose, intended use cases and deployment context of, and benefits afforded by the high-risk artificial intelligence system and (ii) whether the deployment or use of the high-risk artificial intelligence system poses a reasonably foreseeable risk of algorithmic discrimination and, if so, (a) the nature of such algorithmic discrimination and (b) the steps that have been taken, to the extent feasible, to mitigate such risk;

2. For each post-deployment impact assessment completed pursuant to this section, whether the intended use cases of the high-risk artificial intelligence system as updated were consistent with, or varied from, the developer's intended uses of such high-risk artificial intelligence system;

3. A description of (i) the categories of data the high-risk artificial intelligence system processes as inputs and (ii) the outputs such high-risk artificial intelligence system produces;

4. If the deployer used data to customize the high-risk artificial intelligence system, an overview of the categories of data the deployer used to customize such high-risk artificial intelligence system;

5. A description of any transparency measures taken concerning the high-risk artificial intelligence system, including any measures taken to disclose to a consumer in the Commonwealth that such high-risk artificial intelligence system is in use when such high-risk artificial intelligence system is in use; and

6. A description of any post-deployment monitoring performed and user safeguards provided concerning such high-risk artificial intelligence system, including any oversight process established by

the deployer to address issues arising from deployment or use of such high-risk artificial intelligence system as such issues arise.

A single impact assessment may address a comparable set of high-risk artificial intelligence systems deployed or used by a deployer. High-risk artificial intelligence systems that are in conformity with artificial intelligence standards, or parts thereof, shall be presumed to be in conformity with related requirements set out in this section and in associated regulations. If a deployer completes an impact assessment for the purpose of complying with another applicable law or regulation, such impact assessment shall be deemed to satisfy the requirements established in this subsection if such impact assessment is reasonably similar in scope and effect to the impact assessment that would otherwise be completed pursuant to this subsection. A deployer that completes an impact assessment pursuant to this subsection shall maintain such impact assessment and all records concerning such impact assessment for a reasonable period of time.

D. Not later than the time that a deployer uses a high-risk artificial intelligence system to make a consequential decision concerning an individual, the deployer shall notify the individual that the deployer is using a high-risk artificial intelligence system to make such consequential decision concerning such individual and provide to the individual a statement disclosing the purpose of such high-risk artificial intelligence system.

E. Each deployer shall make available, in a manner that is clear and readily available, a statement summarizing how such deployer manages any reasonably foreseeable risk of algorithmic discrimination that may arise from the use or deployment of the high-risk artificial intelligence system.

§ 59.1-606. Exemptions.

A. Nothing in this chapter shall be construed to restrict a developer's or deployer's ability to (i) comply with federal, state, or municipal ordinances or regulations; (ii) comply with a civil, criminal, or regulatory inquiry, investigation, subpoena, or summons by federal, state, municipal, or other governmental authorities; (iii) cooperate with law-enforcement agencies concerning conduct or activity that the developer or deployer reasonably and in good faith believes may violate federal, state, or municipal ordinances or regulations; (iv) investigate, establish, exercise, prepare for, or defend legal claims; (v) provide a product or service specifically requested by a consumer; (vi) perform under a contract to which a consumer is a party, including fulfilling the terms of a written warranty; (vii) take steps at the request of a consumer prior to entering into a contract; (viii) take immediate steps to protect an interest that is essential for the life or physical safety of the consumer or another individual; (ix) prevent, detect, protect against, or respond to security incidents, identity theft, fraud, harassment, malicious or deceptive activities, or any illegal activity, preserve the integrity or security of systems, or investigate, report, or prosecute those responsible for any such action; (x) engage in public or peer-reviewed scientific or statistical research in the public interest that adheres to all other applicable ethics and privacy laws and is approved, monitored, and governed by an institutional review board that determines, or similar independent oversight entities that determine, (a) that the expected benefits of the research outweigh the risks associated with such research and (b) whether the developer or deployer has implemented reasonable safeguards to mitigate the risks associated with such research; (xi) assist another developer or deployer with any of the obligations imposed by this chapter; or (xii) take any action that is in the public interest in the areas of public health, community health, or population health, but solely to the extent that such action is subject to suitable and specific measures to safeguard the public.

B. The obligations imposed on developers or deployers by this chapter shall not restrict a developer's or deployer's ability to (i) conduct internal research to develop, improve, or repair products, services, or technologies; (ii) effectuate a product recall; (iii) identify and repair technical errors that impair existing or intended functionality; or (iv) perform internal operations that are reasonably aligned with the expectations of the consumer or reasonably anticipated based on the consumer's existing relationship with the developer or deployer.

C. The obligations imposed on developers or deployers by this chapter shall not apply where compliance by the developer or deployer with such obligations would violate an evidentiary privilege under the laws of the Commonwealth.

D. Nothing this chapter shall be construed to impose any obligation on a developer or deployer that adversely affects the rights or freedoms of any person, including the rights of any person to freedom of speech or freedom of the press guaranteed in the First Amendment to the Constitution of the United States or under the Virginia Human Rights Act (§ 2.2-3900 et seq.).

E. If a developer or deployer engages in any action pursuant to an exemption set forth in this section, the developer or deployer bears the burden of demonstrating that such action qualifies for such exemption.

§ 59.1-607. Enforcement; civil penalty.

A. The Attorney General shall have exclusive authority to enforce the provisions of this chapter.

B. Whenever the Attorney General has reasonable cause to believe that any person has engaged in or is engaging in any violation of this chapter, the Attorney General is empowered to issue a civil investigative demand. The provisions of § 59.1-9.10 shall apply mutatis mutandis to civil investigative demands issued pursuant to this section.

C. Notwithstanding any contrary provision of law, the Attorney General may cause an action to be brought in the appropriate circuit court in the name of the Commonwealth to enjoin any violation of this chapter. The circuit court having jurisdiction may enjoin such violation notwithstanding the existence of an adequate remedy at law. In any action brought pursuant to this section, it shall not be necessary that damages be proved.

D. Any person who violates the provisions of this chapter shall be subject to a civil penalty in an amount not to exceed $1,000 plus reasonable attorney fees, expenses, and court costs, as determined by the court. Any person who willfully violates the provisions of this chapter shall be subject to a civil penalty in amount not less than $1,000 and not more than $10,000 plus reasonable attorney fees, expenses, and court costs, as determined by the court. Such civil penalties shall be paid into the Literary Fund.

E. Each violation of this chapter shall constitute a separate violation and shall be subject to any civil penalties imposed under this section.

F. The Attorney General may require that a developer disclose to the Attorney General any statement or documentation described in this chapter if such statement or documentation is relevant to an investigation conducted by the Attorney General. The Attorney General may also require that a deployer disclose to the Attorney General any risk management policy designed and implemented, impact assessment completed, or record maintained pursuant to this chapter if such risk management policy, impact assessment, or record is relevant to an investigation conducted by the Attorney General.

2. That the provisions of § 59.1-605 of the Code of Virginia, as created by this act, shall become effective on July 1, 2026.