25106036D

SENATE BILL NO. 1214

AMENDMENT IN THE NATURE OF A SUBSTITUTE

(Proposed by the Senate Committee on General Laws and Technology on January 24, 2025)

(Patron Prior to Substitute—Senator Aird)

A BILL to amend and reenact § 2.2-2007 of the Code of Virginia and to amend the Code of Virginia by adding in Title 2.2 a chapter numbered 55.6, consisting of sections numbered 2.2-5517 through 2.2-5522, relating to high-risk artificial intelligence; development, deployment, and use by public bodies; work group; report.

Be it enacted by the General Assembly of Virginia:

1. That § 2.2-2007 of the Code of Virginia is amended and reenacted and that the Code of Virginia is amended by adding in Title 2.2 a chapter numbered 55.6, consisting of sections numbered 2.2-5517 through 2.2-5522, as follows:

§ 2.2-2007. Powers of the CIO.

A. The CIO shall promulgate regulations necessary or incidental to the performance of duties or execution of powers conferred under this chapter. The CIO shall also develop policies, standards, and guidelines for the planning, budgeting, procurement, development, maintenance, security, and operations of information technology for executive branch agencies. Such policies, standards, and guidelines shall include those necessary to:

1. Support state and local government exchange, acquisition, storage, use, sharing, and distribution of data and related technologies.

2. Support the development of electronic transactions, including the use of electronic signatures as provided in § 59.1-496.

3. Support a unified approach to information technology across the totality of state government, thereby assuring that the citizens and businesses of the Commonwealth receive the greatest possible security, value, and convenience from investments made in technology.

4. Ensure that the costs of information technology systems, products, data, and services are contained through the shared use of existing or planned equipment, data, or services.

5. Provide for the effective management of information technology investments through their entire life cycles, including identification, business case development, selection, procurement, implementation, operation, performance evaluation, and enhancement or retirement. Such policies, standards, and guidelines shall include, at a minimum, the periodic review by the CIO of agency Commonwealth information technology projects.

6. Establish an Information Technology Investment Management Standard based on acceptable technology investment methods to ensure that all executive branch agency technology expenditures are an integral part of the Commonwealth's performance management system, produce value for the agency and the Commonwealth, and are aligned with (i) agency strategic plans, (ii) the Governor's policy objectives, and (iii) the long-term objectives of the Council on Virginia's Future.

B. In addition to other such duties as the Secretary may assign, the CIO shall:

1. Oversee and administer the Virginia Technology Infrastructure Fund created pursuant to § 2.2-2023.

2. Report annually to the Governor, the Secretary, and the Joint Commission on Technology and Science created pursuant to § 30-85 on the use and application of information technology by executive branch agencies to increase economic efficiency, citizen convenience, and public access to state government.

3. Prepare annually a report for submission to the Secretary, the Information Technology Advisory Council, and the Joint Commission on Technology and Science on a prioritized list of Recommended Technology Investment Projects (RTIP Report) based upon major information technology projects submitted for business case approval pursuant to this chapter. As part of the RTIP Report, the CIO shall develop and regularly update a methodology for prioritizing projects based upon the allocation of points to defined criteria. The criteria and their definitions shall be presented in the RTIP Report. For each project recommended for funding in the RTIP Report, the CIO shall indicate the number of points and how they were awarded. For each listed project, the CIO shall also report (i) all projected costs of ongoing operations and maintenance activities of the project for the next three biennia following project implementation; (ii) a justification and description for each project baseline change; and (iii) whether the project fails to incorporate existing standards for the maintenance, exchange, and security of data. This report shall also include trends in current projected information technology spending by executive branch agencies and secretariats, including spending on projects, operations and maintenance, and payments to VITA. Agencies shall provide all project and cost information required to complete the RTIP Report to the CIO prior to May 31 immediately preceding any budget biennium in which the project appears in the Governor's budget bill.

4. Provide oversight for executive branch agency efforts to modernize the planning, development,

implementation, improvement, operations and maintenance, and retirement of Commonwealth information technology, including oversight for the selection, development and management of enterprise information technology.

5. Develop statewide technical and data standards and specifications for information technology and related systems, including (i) the efficient exchange of electronic information and technology, including infrastructure, between the public and private sectors in the Commonwealth and (ii) the utilization of nationally recognized technical and data standards for health information technology systems or software purchased by an executive branch agency.

6. Direct the compilation and maintenance of an inventory of information technology, including but not limited to personnel, facilities, equipment, goods, and contracts for services.

7. Provide for the centralized marketing, provision, leasing, and executing of licensing agreements for electronic access to public information and government services through the Internet, wireless devices, personal digital assistants, kiosks, or other such related media on terms and conditions as may be determined to be in the best interest of the Commonwealth. VITA may fix and collect fees and charges for (i) public information, media, and other incidental services furnished by it to any private individual or entity, notwithstanding the charges set forth in § 2.2-3704, and (ii) such use and services it provides to any executive branch agency or local government. Nothing in this subdivision authorizing VITA to fix and collect fees for providing information services shall be construed to prevent access to the public records of any public body pursuant to the provisions of the Virginia Freedom of Information Act (§ 2.2-3700 et seq.). VITA is authorized, subject to the approval by the Secretary of Administration and any other affected Secretariat, to delegate the powers and responsibilities granted in this subdivision to any agency within the executive branch.

8. Periodically evaluate the feasibility of outsourcing information technology resources and services, and outsource those resources and services that are feasible and beneficial to the Commonwealth.

9. Have the authority to enter into and amend contracts, including contracts with one or more other public bodies, or public agencies or institutions or localities of the several states, of the United States or its territories, or the District of Columbia, for the provision of information technology services.

10. Develop, publish, and maintain policies and procedures concerning the development, procurement, implementation, utilization, and ongoing assessment of systems that employ high-risk artificial intelligence systems, as defined in § 2.2-5517, and are in use by public bodies, consistent with the provisions of Chapter 55.6 (§ 2.2-5517 et seq.). Such policies and procedures shall, at a minimum, (i) govern the procurement, implementation, and ongoing assessment of any such system by a public body; (ii) address and provide resources regarding data security and privacy issues that may arise from the development and deployment of high-risk artificial intelligence systems by public bodies; (iii) be sufficient to ensure that no such system results in any algorithmic discrimination, as defined in § 2.2-5517; (iv) create guidelines for acceptable use policies for public bodies integrating high-risk artificial intelligence systems pursuant to § 2.2-5520; and (v) require a public body to assess the likely impact of any such system before implementing such system and perform ongoing assessments of such system to ensure that no such system results in any such algorithmic discrimination, as defined in § 2.2-5517. Such policies and procedures shall include a requirement that a high-risk artificial intelligence system compliance clause be included in procurement contracts for systems that use a high-risk artificial intelligence system for which negotiation or renegotiation is begun on or after July 1, 2026, requiring compliance with the provisions of Chapter 55.6 (§ 2.2-5517 et seq.) and any other applicable state law governing the development or deployment of high-risk artificial intelligence systems, as applicable.

C. Consistent with § 2.2-2012, the CIO may enter into public-private partnership contracts to finance or implement information technology programs and projects. The CIO may issue a request for information to seek out potential private partners interested in providing programs or projects pursuant to an agreement under this subsection. The compensation for such services shall be computed with reference to and paid from the increased revenue or cost savings attributable to the successful implementation of the program or project for the period specified in the contract. The CIO shall be responsible for reviewing and approving the programs and projects and the terms of contracts for same under this subsection. The CIO shall determine annually the total amount of increased revenue or cost savings attributable to the successful implementation of a program or project under this subsection and such amount shall be deposited in the Virginia Technology Infrastructure Fund created in § 2.2-2023. The CIO is authorized to use moneys deposited in the Fund to pay private partners pursuant to the terms of contracts under this subsection. All moneys in excess of that required to be paid to private partners, as determined by the CIO, shall be reported to the Comptroller and retained in the Fund. The CIO shall prepare an annual report to the Governor, the Secretary, and General Assembly on all contracts under this subsection, describing each information technology program or project, its progress, revenue impact, and such other information as may be relevant.

D. Executive branch agencies shall cooperate with VITA in identifying the development and operational requirements of proposed information technology systems, products, data, and services, including the

proposed use, functionality, and capacity, and the total cost of acquisition, operation, and maintenance.

CHAPTER 55.6.

USE OF HIGH-RISK ARTIFICIAL INTELLIGENCE SYSTEMS.

§ 2.2-5517. Definitions.

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

"Algorithmic discrimination" means any discrimination that results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of their actual or perceived age, color, disability, ethnicity, genetic information, limited proficiency in the English language, national origin, race, religion, reproductive health, sex, sexual orientation, veteran status, or other classification protected under state or federal law. "Algorithmic discrimination" does not include (i) the offer, license, or use of a high-risk artificial intelligence system by a developer, integrator, or deployer for the sole purpose of the developer's, integrator's, or deployer's self-testing to identify, mitigate, or prevent discrimination or otherwise ensure compliance with state and federal law or (ii) the expansion of an applicant, customer, or participant pool to increase diversity or redress historical discrimination.

"Artificial intelligence" means a set of technologies that enables machines to perform tasks under varying and unpredictable circumstances that typically require human oversight or intelligence, or that can learn from experience and improve performance when exposed to data sets.

"Artificial intelligence system" means any machine-based system that, for any explicit or implicit objective, infers from the inputs such system receives how to generate outputs, including content, decisions, predictions, and recommendations, that can influence physical or virtual environments.

"Consequential decision" means any decision that has a material legal, or similarly significant, effect on the provision or denial to any consumer of, or the cost or terms of, (i) education enrollment or an education opportunity, (ii) employment or an employment opportunity, (iii) a financial or lending service, (iv) an essential government service, (v) health care services, (vi) housing, (vii) insurance, or (viii) a legal service.

"Consumer" means a natural person 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 public body that deploys or uses a high-risk artificial intelligence system to make a consequential decision.

"Developer" means any public body 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.

"Facial recognition" means the use of a computer system that, for the purpose of attempting to determine the identity of an unknown individual, uses an algorithm to compare the facial biometric data of an unknown individual derived from a photograph, video, or image to a database of photographs or images and associated facial biometric data in order to identify potential matches to an individual. "Facial recognition" does not include facial verification technology, which involves the process of comparing an image or facial biometric data of a known individual, where such information is provided by that individual, to an image database, or to government documentation containing an image of the known individual, to identify a potential match in pursuit of the individual's identity.

"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.

"General-purpose artificial intelligence model" means any form of artificial intelligence system that (i) displays significant generality, (ii) is capable of competently performing a wide range of distinct tasks, and (iii) can be integrated into a variety of downstream applications or systems. "General-purpose artificial intelligence model" does not include any artificial intelligence model that is used for development, prototyping, or research activities before such artificial intelligence model is released on the market.

"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 substantial 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 sufficient human review, or (iv) perform a preparatory task to an assessment relevant to a consequential decision. There is a rebuttable presumption that "high-risk artificial intelligence system" does not include any of the following technologies:

1. Anti-fraud technology that does not use facial recognition technology;

2. Anti-malware technology;

3. Anti-virus technology;

4. Artificial intelligence-enabled video games;

5. Calculators;
6. Cybersecurity technology;
7. Databases;
8. Data storage;
9. Firewall technology;
10. Internet domain registration;
11. Internet website loading;
12. Networking;
13. Spam and robocall filtering;
14. Spell-checking technology;
15. Spreadsheets;
16. Web caching;
17. Web hosting or any similar technology; or
18. Technology that communicates with consumers in natural language for the purpose of providing users with information, making referrals or recommendations, and answering questions and is subject to an accepted use policy that prohibits generating content that is discriminatory or harmful.

"Integrator" means a public body that knowingly integrates an artificial intelligence system into a software application and places such software application on the market or makes such software application available for public use. An "integrator" does not include a public body offering information technology infrastructure.

"Intentional and substantial modification" means any deliberate change made to (i) an artificial intelligence system that results in any new reasonably foreseeable risk of algorithmic discrimination or (ii) a general-purpose artificial intelligence model that affects compliance of the general-purpose artificial intelligence model, materially changes the purpose of the general-purpose artificial intelligence model, or results in any new reasonably foreseeable risk of algorithmic discrimination. "Intentional and substantial modification" does not include any change made to a high-risk artificial intelligence system, or the performance of a high-risk artificial intelligence system, if (a) the high-risk artificial intelligence system continues to learn after such high-risk artificial intelligence system is offered, sold, leased, licensed, given, or otherwise made available to a deployer, or deployed, and (b) such change (1) is made to such high-risk artificial intelligence system as a result of any learning described in clause (a), and (2) was predetermined by the deployer or the third party contracted by the deployer when such deployer or third party completed the initial impact assessment of such high-risk artificial intelligence system as required in § 2.2-5519.

"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.

"Public body" means any authority, board, department, instrumentality, agency, or other unit of state government. "Public body" does not include any county, city, or town; or any local or regional governmental authority.

"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 and that results in any new or reasonably foreseeable risk of algorithmic discrimination.

"Substantial factor" means a factor that (i) assists in making a consequential decision, (ii) is capable of altering the outcome of a consequential decision, and (iii) is generated by an artificial intelligence system. "Substantial factor" includes any use of an artificial intelligence system to generate any content, decision, prediction, or recommendation concerning a consumer that is used as a basis to make a consequential decision concerning the consumer.

"Synthetic digital content" means any digital content, including any audio, image, text, or video, that is produced or manipulated by a generative artificial intelligence system, including a general-purpose artificial intelligence model.

"Trade secret" means information, including a formula, pattern, compilation, program, device, method, technique, or process, that (i) derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable by proper means by, other persons who can obtain economic value from its disclosure or use and (ii) is the subject of efforts that are reasonable under the circumstances to maintain its secrecy.

§ 2.2-5518. Operating standards for public bodies developing 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:

1. A statement disclosing the intended uses of such high-risk artificial intelligence system;
2. Documentation disclosing the following:
a. The known or reasonably known limitations of such high-risk artificial intelligence system, including

any and all known or 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, sold, given, or otherwise made available to a developer;

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;

3. Documentation describing (i) how the high-risk artificial intelligence system was evaluated for performance and for mitigation of algorithmic discrimination before such system was made available to the deployer; (ii) the data governance measures used to cover the training data sets and the measures used to examine the suitability of data sources, possible biases of data sources, and appropriate mitigation; (iii) the intended outputs of the high-risk artificial intelligence system; (iv) the measures the developer has taken to mitigate known or reasonably foreseeable risks of algorithmic discrimination that may arise from the reasonably foreseeable deployment of the high-risk artificial intelligence system; and (v) how the high-risk artificial intelligence system should be used, not be used, and be monitored by an individual when such system is used to make, or is a substantial factor in making, a consequential decision; and

4. Any additional documentation that is reasonably necessary to assist the deployer in understanding the outputs and monitoring performance of the high-risk artificial intelligence system for risks of algorithmic discrimination.

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 as required in § 2.2-5519.

C. A developer that also serves as a deployer for any high-risk artificial intelligence system shall not be required to generate the documentation required by this section unless such high-risk artificial intelligence system is provided to an unaffiliated entity acting as a deployer or as otherwise required by law.

D. Nothing in this section shall be construed to require a developer to disclose any trade secret.

E. High-risk artificial intelligence systems that are in conformity with the latest version of the Artificial Intelligence Risk Management Framework published by the National Institute of Standards and Technology, Standard ISO/IEC 42001 of the International Organization for Standardization, or another nationally or internationally recognized risk management framework for artificial intelligence systems, or parts thereof, shall be presumed to be in conformity with related requirements set out in this section and in associated regulations.

F. For any disclosure required pursuant to this section, each developer shall, no later than 90 days after the developer performs an intentional and substantial modification to any high-risk artificial intelligence system, update such disclosure as necessary to ensure that such disclosure remains accurate.

§ 2.2-5519. Operating standards for public bodies deploying high-risk artificial intelligence systems.

A. 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, Standard ISO/IEC 42001 of the International Organization for Standardization, 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.

B. 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 subsection, 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 list of any metrics used to evaluate the performance and known limitations of the high-risk artificial intelligence system;

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

7. 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 the latest version of the Artificial Intelligence Risk Management Framework published by the National Institute of Standards and Technology, Standard ISO/IEC 42001 of the International Organization for Standardization, or another nationally or internationally recognized risk management framework for artificial intelligence systems, 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 five years.

C. Not later than the time that a deployer uses a high-risk artificial intelligence system to make a consequential decision concerning a consumer, the deployer shall notify the consumer that the deployer is using a high-risk artificial intelligence system to make such consequential decision concerning such consumer and provide to the consumer a statement disclosing (i) the purpose of such high-risk artificial intelligence system, (ii) the nature of such system, (iii) the nature of the consequential decision, (iv) the contact information for the deployer, and (v) a description in plain language of such system.

If such consequential decision is adverse to such consumer, the deployer shall provide to the consumer (a) a statement disclosing the principal reason or reasons for the consequential decision, including (1) the degree to which and manner in which the high-risk artificial intelligence system contributed to the consequential decision, (2) the type of data that was processed by such system in making the consequential decision, and (3) the sources of such data; (b) an opportunity to correct any incorrect personal data that the high-risk artificial intelligence system processed in making, or as a substantial factor in making, the consequential decision; and (c) an opportunity to appeal such adverse consequential decision concerning the consumer arising from the deployment of such system. Any such appeal shall allow for human review, if technically feasible, unless providing the opportunity for appeal is not in the best interest of the consumer, including instances in which any delay might pose a risk to the life or safety of such consumer.

D. 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.

E. For any disclosure required pursuant to this section, each deployer shall, no later than 90 days after the developer performs an intentional and substantial modification to any high-risk artificial intelligence system, update such disclosure as necessary to ensure that such disclosure remains accurate.

§ 2.2-5520. Operating standards for public bodies integrating high-risk artificial intelligence systems.

Each integrator of a high-risk artificial intelligence system shall develop and adopt an acceptable use

policy, which shall limit the use of the high-risk artificial intelligence system to mitigate known risks of algorithmic discrimination.

Each integrator of a high-risk artificial intelligence system shall provide to the deployer clear, conspicuous notice of (i) the name or other identifier of the high-risk artificial intelligence system integrated into a software application provided to the deployer; (ii) the name and contact information of the developer of the high-risk artificial intelligence system integrated into a software application provided to the deployer; (iii) whether the integrator has adjusted the model weights of the high-risk artificial intelligence system integrated into the software application by exposing it to additional data, a summary of the adjustment process, and how such process and the resulting system were evaluated for risk of algorithmic discrimination; (iv) a summary of any other non-substantial modifications made by the integrator; and (v) the integrator's acceptable use policy.

§ 2.2-5521. Exemptions.

A. Nothing in this chapter shall be construed to restrict a developer's, integrator'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, local, or other governmental authorities; (iii) cooperate with law-enforcement agencies concerning conduct or activity that the developer, integrator, or deployer reasonably and in good faith believes may violate federal, state, or local law, 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, or malicious or deceptive activities; (x) take actions to prevent, detect, protect against, report, or respond to the production, generation, incorporation, or synthesization of child sex abuse material, or any illegal activity, preserve the integrity or security of systems, or investigate, report, or prosecute those responsible for any such action; (xi) 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, integrator, or deployer has implemented reasonable safeguards to mitigate the risks associated with such research; (xii) assist another developer, integrator, or deployer with any of the obligations imposed by this chapter; or (xiii) 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, integrators, or deployers by this chapter shall not restrict a developer's, integrator'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, integrator, or deployer.

C. Nothing in this chapter shall be construed to impose any obligation on a developer, integrator, or deployer to disclose trade secrets.

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

E. Nothing in this chapter shall be construed to impose any obligation on a developer, integrator, or deployer that adversely affects the legally protected 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.).

F. If a developer, integrator, or deployer engages in any action authorized by an exemption set forth in this section, the developer, integrator, or deployer bears the burden of demonstrating that such action qualifies for such exemption.

§ 2.2-5522. Additional requirements.

A. A public body shall not implement any system that employs high-risk artificial intelligence systems unless it has fulfilled the requirements of this section and complied with the provisions of this chapter and the high-risk artificial intelligence policies and procedures developed by the Chief Information Officer of the Commonwealth pursuant to subdivision B 10 of § 2.2-2007.

B. A public body procuring any system that employs high-risk artificial intelligence systems shall in all future contracts for the procurement of such systems for which negotiation or renegotiation is begun on or after July 1, 2026, include a high-risk artificial intelligence system compliance clause, as developed by the

Chief Information Officer of the Commonwealth pursuant to § 2.2-2007.

C. Prior to implementing any system that employs high-risk artificial intelligence systems, the public body shall comply with the impact assessment requirements of § 2.2-5519. A public body shall additionally perform ongoing assessments of such system after implementation. If the public body, or the head of the public body, determines, in its discretion, that such system does not comply with such requirements, the public body shall not implement such system or shall cease to use such system to the extent such system does not comply with such requirements.

D. All public bodies that implement high-risk artificial intelligence systems shall annually report on initial and ongoing system assessments and provide an inventory of such systems used. Public bodies in the legislative branch shall submit such report and inventory to the General Assembly. Public bodies in the judicial branch shall submit such report and inventory to the Executive Secretary of the Supreme Court of Virginia. Public bodies in the executive branch and any other public bodies not specified in this subsection shall submit such report and inventory to the Chief Information Officer of the Commonwealth. Such report and inventory shall be transmitted to the appropriate entity annually.

2. That the Chief Information Officer of the Commonwealth (CIO) shall convene a work group to examine the impact on and the ability of local governments to comply with the requirements of this act. The work group shall consist of a representative from the Virginia Association of Counties who is also a representative of a member county, a representative from the Virginia Municipal League who is also a representative of a member locality, a representative of the Virginia Association of Chiefs of Police, a representative from the Virginia Association of Commonwealth's Attorneys, the chief information officer of a school division, the chief information officer of a county, the chief information officer of a city, a representative from the Department of Human Resource Management, a representative of a regional technology council, a member of the Joint Commission on Technology and Science (JCOTS) who is a member of the House of Delegates, and a member of JCOTS who is a member of the Senate. The CIO shall submit a report of the work group's findings to JCOTS no later than December 1, 2025.

3. That the provisions of the first enactment of this act shall become effective on July 1, 2026.